


Strong-field gravitational lensing by black holes

Jake O. Shipley



A thesis submitted for the degree of
Doctor of Philosophy

School of Mathematics and Statistics
University of Sheffield

September 2019

Summary

In this thesis we study aspects of strong-field gravitational lensing by black holes in general relativity, with a particular focus on the role of integrability and chaos in geodesic motion.

We begin with a review of theoretical aspects of Einstein’s theory of general relativity, as well as important concepts and techniques from the fields of non-linear dynamics and chaos theory. Next, we review the topic of gravitational lensing. The history of gravitational lensing and the perturbative lensing formalism are discussed. We then present an overview of gravitational lensing from a spacetime perspective, including the leading-order geometric optics approximation for high-frequency electromagnetic waves on curved spacetime; the role of (unstable) photon orbits; and black hole shadows.

We investigate binary black hole shadows using the Majumdar–Papapetrou static binary black hole (or di-hole) solution. We demonstrate that the propagation of null geodesics on this spacetime background is a natural example of chaotic scattering. The role of unstable photon orbits is discussed. We develop a symbolic dynamics to describe null geodesics and to understand the structure of one-dimensional binary black hole shadows. We demonstrate that, in situations where chaotic scattering is permitted, the shadows exhibit a self-similar fractal structure akin to the Cantor set. A gallery of two-dimensional binary black hole shadows, realised using backwards ray-tracing, is presented and analysed in detail.

Next, we use techniques from the field of non-linear dynamics to quantify fractal structures in the shadows of binary black holes, using the static Majumdar–Papapetrou di-hole as a toy model. Our perspective is that binary black hole shadows can be viewed as the exit basins of an open Hamiltonian system with three escapes. We compute the uncertainty exponent – a quantity related to the fractal dimension – for one-dimensional binary black hole shadows. Using a recently developed numerical algorithm, called the merging method, we demonstrate that parts of the Majumdar–Papapetrou di-hole shadow may possess the Wada property: any point on the boundary of one basin is on the boundary of at least two additional basins.

We study the existence, stability and phenomenology of circular photon orbits in stationary axisymmetric four-dimensional spacetimes in general relativity. We use a Hamiltonian formalism to describe null geodesics of the Weyl–Lewis–Papapetrou geometry. Using the Einstein–Maxwell equations, we demonstrate that generic stable photon orbits are forbidden in pure vacuum, but may arise in electrovacuum. As a case study, we consider stable photon orbits around the Reissner–Nordström family of static di-holes. We examine the onset of chaos in the motion of bounded null geodesics using Poincaré sections.

In the final chapter, we apply a higher-order geometric optics formalism to describe the propagation of electromagnetic waves on Kerr spacetime. Our principal motivation is to calculate the sub-dominant correction to the electromagnetic stress–energy tensor. We use a complex self-dual bivector, built from the closed conformal Killing–Yano tensor and its Hodge dual, to construct a complex null tetrad which is parallel-propagated along null geodesics. We introduce a system of transport equations to calculate certain Newman–Penrose and higher-order geometric optics quantities. We derive generalised power series solutions to these transport equations through sub-leading order in the neighbourhood of caustic points. Finally, we introduce a practical method which may be used to evolve transport equations for divergent quantities through caustic points.

Acknowledgements

I would like express my sincere gratitude to Dr. Sam Dolan for his support, guidance and encouragement throughout the last four years. I could not have wished for a more enthusiastic, knowledgeable and patient supervisor; it has been a pleasure and a privilege to learn from him and work with him. I would especially like to thank Sam for going above and beyond the role of supervisor. No matter how busy he is, his door is always open; he is a constant source of advice and inspiration; and he is always the first to buy a round at the pub! I am extremely fortunate to be able to call Sam my friend, as well as my supervisor.

Chapter 5 of this thesis is the result of a thoroughly enjoyable collaboration with Professor Miguel Sanjuán and Dr. Álvaro Daza. I thank them both for their warm welcome and generous hospitality during my visit to Madrid in December 2016, and for teaching me with such enthusiasm about non-linear dynamics and chaos theory.

I acknowledge financial support from the University of Sheffield's Harry Worthington Scholarship. This funding has allowed me to pursue my research interests for four years, and it has provided me with the opportunity to visit a host of amazing places – including Paris, Aveiro, Dublin, Madrid and Rome – all in the name of research.

It has been a pleasure to be part of the Cosmology, Relativity and Gravitation research group for the last four years. I would like to acknowledge the kindness and friendship of all members past and present. Special thanks go to Carl Kent, Jack Morrice, David Dempsey, Tom Morley, Tom Stratton, Jake Percival, Vis Balakumar, Elisa Maggio, Mohamed Ould El Hadj, Carsten van de Bruck and Elizabeth Winstanley.

Lastly, I would like to thank my parents, sisters and Hannah for their encouragement, support and unconditional love throughout the last four years and beyond. It would take an acknowledgements section the size of this thesis to even begin to thank you all properly.

Contents

1	Introduction	1
2	Dynamics in general relativity	10
2.1	Geometry of spacetime	10
2.2	Geodesics	21
2.3	Field equations and solutions	26
2.4	Black holes	29
2.5	Newman–Penrose formalism	32
2.6	Integrability and chaos in general relativity	37
3	Gravitational lensing	53
3.1	Introduction	53
3.2	Perturbative lensing	54
3.3	Non-perturbative lensing	59
4	Binary black hole shadows and chaotic scattering	70
4.1	Introduction	70
4.2	Majumdar–Papapetrou solution	72
4.3	Null geodesic structure	73
4.4	Planar geodesics and one-dimensional shadows	76
4.5	Symbolic dynamics	78
4.6	Cantor-like structure of the one-dimensional shadow	82
4.7	Non-planar geodesics and two-dimensional shadows	93
4.8	Following rays through the event horizons	105
4.9	Discussion	109
5	Fractal structures in binary black hole shadows	113
5.1	Introduction	113
5.2	Majumdar–Papapetrou open Hamiltonian system	114
5.3	Transition in shadow structure	119
5.4	Uncertainty exponent and fractal dimension	122
5.5	The Wada property	128
5.6	Numerical test for Wada basins	130
5.7	Discussion	137

6	Stable photon orbits in stationary axisymmetric spacetimes	140
6.1	Introduction	140
6.2	Photon orbits in Kerr–Newman spacetime	141
6.3	Photon orbits in stationary axisymmetric spacetimes	145
6.4	Existence of stable photon orbits around di-holes	148
6.5	Geodesic structure of stable photon orbits	153
6.6	Discussion	158
7	Higher-order geometric optics on Kerr spacetime	162
7.1	Introduction	162
7.2	Higher-order geometric optics formalism	164
7.3	Kerr spacetime	166
7.4	Construction of complex null tetrads	171
7.5	Newman–Penrose formalism and transport equations	176
7.6	Far-field asymptotics	180
7.7	Wavefronts, caustic points and transport equations	183
7.8	Evolving transport equations through caustic points	190
7.9	Discussion	201
	Appendices	205
A	Translation between symbolic codes	206
B	Circular photon orbits of the Majumdar–Papapetrou di-hole	208
C	Hénon–Heiles Hamiltonian system	212
D	Bretón–Manko–Aguilar di-hole solution	220
E	Tetrad calculations	222
	Bibliography	250

Chapter 1

Introduction

On 10 April 2019, the Event Horizon Telescope (EHT) collaboration reported the first image of a black hole [1–6]. Using very-long-baseline interferometry (VLBI), the EHT – a global network of telescopes observing at millimetre wavelengths – captured exquisite-resolution images of radio emission from the supermassive black hole candidate M87*, which is believed to lie at the heart of the giant elliptical galaxy Messier 87 (M87). A key feature of the EHT’s images is a bright asymmetric ring which surrounds a dark central region – the *black hole shadow*. (See Figure 3 of [1], for example.) Overall, the observed images are consistent with expectations for a rotating Kerr black hole, as predicted by Einstein’s general theory of relativity. This ground-breaking observation confirms the existence of black holes – a key prediction of general relativity – and provides a new way to test Einstein’s theory in its most extreme limit.

The observation of M87* is in fact a detection of a *gravitational lensing* effect, i.e., the deflection of light by gravity [7]. According to general relativity, black holes (and all other massive bodies) generate spacetime curvature, which leads to the deviation in the paths of photons as they trace out null geodesics on curved spacetime. The outline of the black hole shadow observed by the EHT is associated with the black hole’s unstable *light-ring*, where spacetime curvature is so strong that light is able to orbit the black hole. Beyond the light-ring, radially infalling photons are doomed to plunge into the black hole, crossing its *event horizon* – a one-way causal boundary in spacetime beyond which nothing can escape to infinity.

The shadow observed by the Event Horizon Telescope encodes important information about the black hole and the spacetime geometry close to the event horizon. Comparing the observed shadow of M87* with a library of ray-traced general-relativistic magnetohydrodynamic simulations, the EHT collaboration have been able to infer a mass of $M = (6.5 \pm 0.7) \times 10^9 M_\odot$ [1, 6]. Moreover, the observations are consistent with the black hole’s spin axis being oriented at 17 degrees from the line of sight, with the black hole rotating in the clockwise direction (i.e., the spin axis points away from us) [1]. It is hoped that future high-resolution VLBI imaging of M87* and other supermassive black hole candidates will allow scientists to test the spin and inclination of the black holes [8], and to continue to probe general relativity in the strong-field regime [9–13].

In 2015, the Laser Interferometer Gravitational-Wave Observatory (LIGO) collabora-

tion detected a gravitational-wave signal emitted by a pair of merging stellar-mass black holes [14]. The characteristic “chirp” profile of the gravitational-wave signal, dubbed GW150914, is consistent with the inspiral, merger and ringdown phases of waveform templates generated by numerical relativity simulations. As well as providing compelling evidence for the existence of both gravitational waves and binary black holes in nature, this detection signalled the birth of gravitational-wave astronomy as an observational science.

Two years later, in 2017, a gravitational-wave signal from a pair of coalescing neutron stars (GW170817) was observed by the LIGO–Virgo collaboration [15]. This signal was accompanied, approximately 1.7 seconds later, by a gamma-ray burst; and its optical transient was detected via a host of observations across the electromagnetic spectrum [16]. This was evidence of another important prediction of Einstein’s theory of general relativity: gravitational waves propagate outwards from their source at the speed of light. Moreover, GW170817 was the first gravitational-wave signal which was observed alongside an electromagnetic counterpart, marking an exciting breakthrough in the field of multi-messenger astronomy.

In total, LIGO–Virgo observed ten gravitational-wave signals during its first two observing runs, of which nine are consistent with the gravitational-wave signal generated by merging black holes, and the other with a binary neutron star merger [17, 18]. The LIGO–Virgo collaboration began their third observing run on 1 April 2019, and a host of candidate gravitational-wave signals have since been detected from compact binary coalescences. With the improvements made to the LIGO–Virgo detectors for the third observing run, there is an exciting possibility that astronomers may detect several binary neutron star mergers, in addition to one or more neutron star–black hole mergers [19].

The ground-breaking discoveries of gravitational waves and the observation of M87* have reignited interest in Einstein’s theory of relativity and gravitational physics. With further ground- and space-based gravitational-wave missions (e.g. KAGRA [20], LISA [21] and others [22]) on the horizon, and images of Sagittarius A* – the supermassive black hole at the centre of the Milky Way – expected soon [23], we are on the cusp of an exciting new era of gravitational astronomy and high-precision experimental tests of extreme gravity.

Whilst astrophysical black holes have only recently been observed directly (in both electromagnetic- and gravitational-wave channels), the theoretical properties of black holes have been studied for decades.

In 1916, just one year after the publication of Einstein’s general theory of relativity, Schwarzschild [24] found an exact solution to the vacuum field equations which describes the spacetime geometry of a static spherically symmetric gravitational source. In fact, Schwarzschild’s solution describes a black hole, although it took a range of detailed analyses performed over a number of years to build up a more complete understanding of the solution’s properties. A generalisation of the Schwarzschild solution, in which the gravitational source is endowed with an electric charge, was revealed independently by Reissner [25] and Nordström [26] between 1916 and 1918. The Reissner–Nordström solution is an exact solution to the Einstein–Maxwell field equations of general relativity and electromagnetism which exhibits the same spacetime symmetries as the uncharged Schwarzschild solution. In 1923, Birkhoff [27] proved that the Schwarzschild spacetime is

the unique spherically symmetric solution of the Einstein field equations in pure vacuum, a result now known as *Birkhoff's theorem*. (In fact, Birkhoff's theorem was first published in 1921 by Jebsen [28].) High levels of symmetry also play an important role in geodesic motion on Schwarzschild spacetime: in spherical coordinates, which are well-adapted to the spacetime symmetries, geodesic motion is separable, and therefore integrable (in the sense of Liouville) [29].

Interest in the properties of very compact gravitational sources continued throughout the 1930s. Key contributions included Chandrasekhar's work on white dwarf stars, and the work of Landau [30] and others [31–33] on neutron stars. At the end of the decade, Oppenheimer and Snyder [34] described the gravitational collapse of massive stars, and the subsequent formation of black holes.

Over the next few decades, the field of black hole physics went through a period of relative dormancy, until Kerr [35] discovered an exact solution to Einstein's field equations describing a rotating black hole in vacuum in 1963. This marked the beginning of the so-called golden age of black hole physics, and a range of influential perspectives on the theoretical properties of black holes followed. These included the generalisation of Kerr's solution (to include electric charge) by Newman and collaborators [36,37]; the formulation and proof of uniqueness and “no-hair” theorems for Kerr black holes [38–40]; the proposal of the Penrose process, a mechanism which allows for energy extraction from rotating black holes [41]; an improved understanding of singularities in general relativity, and the proposal of the cosmic censorship conjecture [41,42]; the formulation of the laws of black hole mechanics [43] and a development of the concept of black hole entropy [44]; and the prediction of Hawking radiation and black hole evaporation [45].

In Boyer–Lindquist coordinates [46], the stationarity and axial symmetry of the spacetime are manifest: the corresponding isometries are encoded in the existence of a pair of commuting Killing vectors. However, it is not immediately obvious that any further symmetries exist. In 1968, Carter [47] demonstrated that there exists a “hidden” symmetry of Kerr spacetime – a fourth integral of motion which permits the separability of the geodesic equations on the spacetime background. This conserved quantity is known as the *Carter constant*. Later, Floyd [48] and Penrose [49] demonstrated that the Kerr spacetime admits a Killing–Yano tensor, whose “square” is the Killing tensor. Subsequent investigations demonstrated that there is a deep geometrical reason behind the existence of the Carter constant: the Kerr spacetime admits a rank-two closed conformal Killing–Yano tensor (or *principal tensor*), which gives rise to a family of Killing tensors. These Killing tensors may then be used to generate the full set of explicit and hidden symmetries on Kerr spacetime. For a comprehensive overview of this topic, see the review by Frolov *et al.* [50] and references therein.

The existence of the principal tensor and its associated “hidden” symmetry is implicated in a number of key results. For example, parallel-transport along geodesics is straightforward [51,52]; gravitational Faraday rotation of the polarisation plane of light is trivial [49]; the Hamilton–Jacobi, Schrödinger [47], Klein–Gordon [53], Dirac [54,55], Rarita–Schwinger [56] and Proca [57] equations are all separable in Boyer–Lindquist coordinates; and the Maxwell equations [58] and those governing gravitational perturba-

tions [59] are separable for the Maxwell and Weyl scalars of extreme spin weight in a certain complex null tetrad.

On Kerr spacetime, the high levels of symmetry ensure the integrability of geodesic motion, which means that trajectories are highly “ordered” in phase space. Moreover, the existence of the principal tensor underpins a range of important results relating to particle motion and wave propagation on Kerr spacetime (as described above). In general stationary axisymmetric spacetimes, however, such results do not hold. On one hand, this renders the study of particle motion and wave propagation a more technically demanding task; on the other, the lack of symmetry gives rise to the possibility of rich chaotic motion, which is typically associated with a range of distinct phenomena, particularly in the field of strong gravitational lensing by black holes and other (ultra-)compact objects [60].

As well as marking the announcement of the first direct detection of a black hole using electromagnetic telescopes by the EHT [1], the year 2019 was the centenary of the *first* detection of a gravitational lensing event. On 29 May 1919, two expeditions – one to Príncipe, an island of the west coast of Africa, and the other to Sobral, a city in north-east Brazil – carried out experiments during a total solar eclipse to measure the deflection of starlight by our own Sun [61]. (For a recent review, see the article by Crispino and Kennefick [62].) The detection of gravitational light deflection provided experimental verification of Einstein’s then novel theory of gravitation and was influential evidence for its superiority over the Newtonian theory [63]. Since then, the field of gravitational lensing has undergone a series of interesting developments – both theoretically and observationally.

In the two decades that followed the 1919 detection, various gravitational lensing phenomena – such as multiple images and Einstein rings – were proposed by a range of authors [64, 65]. In the 1960s, interest in the field was rekindled after the development of the quasi-Newtonian (or *perturbative*) lensing formalism by Refsdal [66] and others [67, 68]. This was followed, in 1978, by the discovery of the multiply imaged quasar Q0957+561 [69] – the first experimental detection of a gravitational lensing effect since the observation of light deflection by the Sun almost sixty years earlier. To date, a host of gravitational lensing phenomena have been observed, including multiply imaged sources, Einstein rings, giant luminous arcs, image distortion, and galactic microlensing; for a review, see [70].

In many cases of interest, gravitational lensing effects are well-described by the perturbative lensing formalism, which is based on a first-order post-Newtonian approximation to general relativity [71]. For example, even in the formation of spectacular lensing phenomena, such as giant luminous arcs, photons are deflected by no more than a few arcseconds [70]. However, in many other situations (e.g. in the strong-field regions around black holes or other compact objects), the quasi-Newtonian formalism breaks down, and a more careful treatment is required. In order to adequately describe the strong-field lensing effects associated with extreme compact objects, one must employ *non-perturbative lensing* (also known as *lensing from a spacetime perspective*), using the full theory of general relativity [7]. In this approach, photons propagate along the null geodesics of a (four-dimensional) Lorentzian spacetime, which is typically assumed to be a solution to Einstein’s field equations of general relativity (or some alternative theory of gravity).

The non-perturbative lensing formalism is particularly well-equipped to deal with the

study of strong-field gravitational lensing by black holes. Interest in the theoretical aspects of gravitational lensing by black holes dates back to the 1960s. In a seminal paper, Synge [72] calculated the angular radius of the shadow of a Schwarzschild black hole, as seen by a distant observer. Then, in 1979, Luminet [73] determined the apparent optical image of a Schwarzschild black hole illuminated by a distant light source, as well as the astrophysically interesting case of the image of a Schwarzschild black hole surrounded by an emitting accretion disk. The simulated images obtained by Luminet bear a striking resemblance to the (real) black hole shadow images captured by the EHT [1].

Since the early work of Synge and Luminet, there has been much interest in the analysis of the strong-field gravitational lensing effects of black holes (and other compact objects). For example, theorists have built up an understanding of the geodesic dynamics on black hole spacetimes [74–76]; explored the existence, stability and phenomenology of light-rings [77–79]; and investigated the structure of black hole shadows [9, 60, 80]. Taking into account the continuing efforts to deepen our understanding of the theoretical aspects of strong-field gravitational lensing, and an ever-increasing ability to test the theoretical predictions of general relativity at exquisite levels of precision, it is clear that gravitational lensing will either play a key role confirming Einstein’s theory as a fundamental law of nature, or perhaps open the door to exciting new physics.

Outline

In this thesis, we analyse theoretical aspects of gravitational lensing by black holes in general relativity, with a particular focus on the role of (non-)integrability, order and chaos. Chapters 2 and 3 are devoted to a review of key concepts and techniques from the fields of general relativity, dynamical systems, chaos theory and gravitational lensing. The remaining chapters present new work carried out by the author and collaborators on the broad theme of strong-field gravitational lensing by black holes. In Chapters 4 and 5, we investigate chaotic scattering and fractal structures in binary black hole shadows. Chapter 6 studies the existence and phenomenology of stable photon orbits in stationary axisymmetric spacetimes. Finally, in Chapter 7, we study the propagation of electromagnetic radiation on Kerr spacetime by applying an extended geometric optics formalism. We conclude the present chapter with a more detailed chapter-by-chapter account of the work presented in this thesis.

Chapter 2. Dynamics in general relativity

In Chapter 2, we review a range of important mathematical tools required for the study of general relativity. In particular, we introduce some aspects of the theory of differential geometry, including differentiable manifolds; vectors, one-forms and tensors; the metric tensor; covariant differentiation and parallel transport; Lie differentiation; Killing vectors; stationarity and staticity; and curvature. We then proceed to look at geodesics, presenting the Lagrangian and Hamiltonian formalisms for geodesic motion on curved spacetime, and deriving the geodesic deviation equation. Next, we present Einstein’s field equations of general relativity, and the Einstein–Maxwell equations of gravity and electromagnetism.

Key solutions are reviewed. We then discuss black holes and introduce the Schwarzschild, Kerr and Kerr–Newman geometries. We give a brief overview of the tetrad formalism, before looking at a special case – the Newman–Penrose formalism – in more detail. We conclude the chapter with a discussion of themes which are central to this thesis: integrability and chaos in dynamical systems, with an emphasis on geodesic motion in general relativity.

Chapter 3. Gravitational lensing

Chapter 3 is devoted to a review of gravitational lensing, which accounts for all effects of a gravitational field on the propagation of electromagnetic radiation. We divide the field into two subfields: perturbative lensing, which uses quasi-Newtonian approximations to describe weak gravitational fields; and non-perturbative lensing (or gravitational lensing from a spacetime perspective), in which light propagates along null geodesics in a four-dimensional Lorentzian spacetime which is a solution to the field equations of general relativity. We begin with a discussion of perturbative lensing. This includes a review of the history of the field as an observational science, beginning with the detection of gravitational light deflection by the Sun in 1919; a discussion of gravitational lensing phenomena which have been observed to date; and a brief account of the mathematical formalism of perturbative lensing, including an illustration of how this can be used to calculate the deflection angle of light due to a static quasi-Newtonian gravitational field. We then proceed to a discussion of gravitational lensing from a spacetime perspective. The theory of electromagnetism in curved spacetime is introduced in a fully covariant manner. We then review the leading-order geometric optics approximation for electromagnetism, which relies on the fundamental assumption that the wavelength (and inverse frequency) is significantly shorter than all other characteristic length (and time) scales, such as the scale set by the spacetime curvature. This scheme reduces the problem of solving wave equations (i.e., Maxwell’s equations) on curved spacetime to one of solving transport equations along the rays (i.e., null geodesics) of the geometry. The application of geometric optics to the study of gravitational lensing phenomena is discussed. We conclude the chapter with an introduction to strong-field gravitational lensing effects associated with black holes, including (unstable) photon orbits and black hole shadows.

Chapter 4. Binary black hole shadows and chaotic scattering

In Chapter 4, we investigate the qualitative features of binary black hole shadows using the Majumdar–Papapetrou binary black hole (or *di-hole*) solution to the Einstein–Maxwell equations, which describes a pair of extremally charged black holes in static equilibrium. We advance the view that the propagation of null geodesics on a binary black hole spacetime is a natural example of chaotic scattering. We find that the existence of two or more dynamically connected fundamental photon orbits gives rise to an uncountable infinity of non-escaping null orbits (comprising the countable set of periodic orbits and the uncountable set of aperiodic orbits), which generate scattering singularities in the initial data. Using the Gaspard–Rice three-disc scatterer as a guide, we develop an appropriate

symbolic dynamics to describe null geodesics. We compare and contrast our approach – referred to here as *decision dynamics* – with an existing symbolic dynamics for the Majumdar–Papapetrou di-hole, which we refer to as *collision dynamics*. We then demonstrate that our symbolic dynamics may be used to construct a one-dimensional binary black hole shadow on initial data, using an iterative procedure akin to the construction of the Cantor set; this argument demonstrates that the one-dimensional binary black hole shadow is self-similar. We then proceed by analysing non-planar null geodesics, aiming to quantify the effect of varying the (conserved) azimuthal angular momentum on the existence and properties of the fundamental null orbits. Using the Hamiltonian formalism for null geodesics, we introduce an effective potential which is independent of the photon’s orbital parameters (i.e., energy and angular momentum); we use this effective potential to understand and classify the null geodesics of the Majumdar–Papapetrou di-hole geometry. In our analysis, we uncover an unexpected feature: the existence of stable bounded photon orbits around Majumdar–Papapetrou di-holes separated by dimensionless coordinate distance $\frac{d}{M} \in \left(\sqrt{\frac{16}{27}}, \sqrt{\frac{32}{27}} \right)$. (This is explored in greater detail in Chapter 6.) We then use ray-tracing to generate two-dimensional binary black hole shadow images for the equal-mass Majumdar–Papapetrou di-hole. We find that a two-dimensional binary black hole shadow image can be decomposed into one-dimensional shadow images which can exhibit regular or Cantor-like structure. We propose a method for tracking null rays through the event horizons of the two black holes in the maximally extended Majumdar–Papapetrou di-hole spacetime. We conclude this chapter with a discussion of related work on black hole shadows and chaotic gravitational lensing signatures. An algorithm for passing between the symbolic codes discussed in this chapter is given in Appendix A. Supplementary calculations on the existence of circular photon orbits around the Majumdar–Papapetrou di-hole can be found in Appendix B.

Chapter 5. Fractal structures in binary black hole shadows

Chapter 5 extends on the work presented in Chapter 4. In particular, we employ techniques from the field of non-linear dynamics to characterise the fractal structures which arise in the shadows of the Majumdar–Papapetrou binary black hole system. We review the scattering of null geodesics in the Majumdar–Papapetrou di-hole spacetime from the perspective of Hamiltonian dynamics. We construct the exit basins in phase space, and highlight qualitative similarities between the Majumdar–Papapetrou system and the Hénon–Heiles Hamiltonian system, a paradigmatic model of two-dimensional time-independent chaotic scattering in Hamiltonian mechanics. (A pedagogical review of the Hénon–Heiles system is given in Appendix C.) We also discuss the structure of black hole shadows, which are viewed as exit basin diagrams on the image plane of a distant observer. We review the uncertainty exponent and present a numerical method to calculate this quantity. We test and calibrate our method by applying it to a simple model – the Cantor basins – for which exact results are known. We then calculate the uncertainty exponent numerically for Majumdar–Papapetrou di-hole shadows. This successfully differentiates between fractal and regular (i.e., non-fractal) regions of the black hole shadow, and agrees with the theoretical predictions of Chapter 4. Next, we apply a recently developed algorithm – the

merging method – to demonstrate that parts of the Majumdar–Papapetrou di-hole shadow possess the Wada property: any point on the boundary of one basin is on the boundary of at least two additional basins. The algorithm is able to successfully distinguish between the Wada and non-Wada parts of the binary black hole shadow.

Chapter 6. Stable photon orbits in stationary axisymmetric spacetimes

In Chapter 6, we explore the existence and phenomenology of *stable* photon orbits in four-dimensional stationary axisymmetric electrovacuum solutions to the Einstein–Maxwell equations. We review the Kerr–Newman solution, and give an overview of the classification of its equatorial circular photon orbits in the charge–spin parameter space. Using a Hamiltonian formalism for rays, we demonstrate that the null geodesics of the stationary axisymmetric Weyl–Lewis–Papapetrou spacetime in four dimensions may be understood by introducing a pair of two-dimensional effective potentials. The fixed points of these potentials correspond to the circular null orbits of the geometry. Restricting attention to the electrovacuum case, we employ a subset of the Einstein–Maxwell equations to classify the fixed points of the effective potentials. We arrive at the following key result for four-dimensional stationary axisymmetric spacetimes: generic stable photon orbits are not permitted in pure vacuum, but may arise in electrovacuum. We investigate the existence and stability of photon orbits around Reissner–Nordström static di-holes, a two-parameter subfamily of the general Bretón–Manko–Aguilar di-hole (reviewed in Appendix D). The Reissner–Nordström static di-hole family includes the uncharged Weyl–Bach di-hole and the extremal Majumdar–Papapetrou di-hole as special cases. In scenarios with high levels of symmetry (e.g. the equal-mass Majumdar–Papapetrou di-hole), we derive closed-form expressions for stable photon orbit existence regions in parameter space. In cases with less symmetry, we employ a numerical method to search for stable photon orbits in parameter space. Finally, using Poincaré sections, we explore the transition from order to chaos for null rays which are bounded in a toroidal region around the black holes in the Majumdar–Papapetrou di-hole geometry. Intriguingly, we find that the Poincaré sections and bounded trajectories of the equal-mass Majumdar–Papapetrou di-hole with dimensionless coordinate separation parameter $\frac{d}{M} = 1$ share many qualitative features with those of the Hénon–Heiles Hamiltonian system (cf. Chapter 5 and Appendix C for the case of *unbounded* trajectories).

Chapter 7. Higher-order geometric optics on Kerr spacetime

In Chapter 7, we apply an extended geometric optics formalism to understand aspects of strong-field gravitational lensing on Kerr spacetime, with the principal aim of computing the sub-dominant correction to the electromagnetic stress–energy tensor. We begin with a review of leading-order geometric optics for the electromagnetic field on an arbitrary curved spacetime, before reviewing the higher-order extension to the geometric optics formalism, originally presented by Dolan [81]. We discuss the geometry of Kerr spacetime, with an emphasis on the closed conformal Killing–Yano tensor, related Killing objects, explicit and hidden symmetries on spacetime and in phase space, and the implications

for the separability and complete integrability of null geodesic motion. We construct a complex null tetrad using a self-dual bivector built from the closed conformal Killing–Yano tensor and its Hodge dual. Performing a Lorentz transformation, we are able to transform this tetrad to a new one which is parallel-transported along null geodesics. We discuss the Newman–Penrose formalism from the perspective of (higher-order) geometric optics on a Ricci-flat spacetime. In particular, we present a closed system of transport equations for the complex Newman–Penrose scalars, and a system of transport equations for certain directional derivatives of Newman–Penrose scalars. We calculate the complex Weyl curvature scalars and their directional derivatives in the parallel-transported complex null tetrad. We determine the far-field behaviour of the Weyl scalars, Newman–Penrose scalars and higher-order geometric optics quantities as generalised power series in r , the radial Boyer–Lindquist coordinate. We discuss wavefronts in geometric optics, and describe caustics, where neighbouring rays cross. The transport equations for Newman–Penrose quantities break down at caustic points. We derive near-caustic solutions to these transport equations as generalised power series in the affine parameter through sub-leading order. Finally, we present a practical method which may be employed to evolve transport equations for divergent quantities through caustic points. We comment on this method and its numerical implementation. Some explicit calculations for this chapter are contained in Appendix E.

Notation and conventions

In this thesis, we adopt the sign conventions of Misner, Thorne and Wheeler [29]. The four-dimensional spacetime metric has Lorentzian signature $(-, +, +, +)$. The Riemann tensor and Ricci tensor are defined in Chapter 2. The Einstein summation convention for repeated indices is assumed throughout. Spacetime indices are denoted by Latin letters from the beginning of the alphabet (e.g. a, b, \dots); spatial indices are denoted by Latin letters from the middle of the alphabet (e.g. i, j, \dots). We employ geometrised units in which the speed of light c and the gravitational constant G are set to unity. Occasionally, we reinsert dimensional constants to aid physical interpretation. All other notation and conventions will be introduced as required.

Chapter 2

Dynamics in general relativity

2.1 Geometry of spacetime

General relativity describes space and time as a four-dimensional continuum known as *spacetime*. The mathematical machinery required to describe spacetime is the differential geometry of pseudo-Riemannian manifolds. This section is devoted to a review of elements of this theory which are important for the study of general relativity. A more exhaustive treatment of topics covered in this section can be found in [29, 82, 83], for example.

2.1.1 Differentiable manifolds

Intuitively, an n -dimensional manifold \mathcal{M} is a space which looks locally like Euclidean space \mathbb{R}^n , but which may have different global properties. Before formally defining a manifold, it will be necessary to introduce some preliminary definitions and terminology.

The *open ball* in \mathbb{R}^n of radius r centred on the point p , denoted $b_r(p)$, is the collection of points $x \in \mathbb{R}^n$ such that $|x - p| < r$, where $|\cdot|$ denotes the Euclidean norm on \mathbb{R}^n .

A subset $U \subset \mathbb{R}^n$ is called *open* if it can be expressed as a union of open balls. Equivalently, the set U is open if, for every point $p \in U$, there exists some real number $\varepsilon > 0$ such that $b_\varepsilon(p) \subset U$.

A *manifold* is a set \mathcal{M} such that any point $p \in \mathcal{M}$ has a neighbourhood $\mathcal{U} \subset \mathcal{M}$ which is homeomorphic to the interior of the n -dimensional unit ball. In order to give a more precise mathematical definition of a manifold, we first require some additional terminology.

On a manifold \mathcal{M} , a *chart* (\mathcal{U}, ψ) consists of a subset $\mathcal{U} \subset \mathcal{M}$ (called a *chart neighbourhood*) and a bijection $\psi: \mathcal{U} \rightarrow U \subseteq \mathbb{R}^n$ (called a *chart map*). The chart map ψ assigns to each point $p \in \mathcal{M}$ an n -tuple of real variables (x^1, \dots, x^n) , which are called *local coordinates*.

Suppose that (\mathcal{U}, ψ) and (\mathcal{U}', ψ') are two charts for \mathcal{M} such that $\mathcal{U} \cap \mathcal{U}' \neq \emptyset$. The *transition map* from $\psi(\mathcal{U} \cap \mathcal{U}') \subset U$ to $\psi'(\mathcal{U} \cap \mathcal{U}') \subset U'$ is the map defined by $\psi' \circ \psi^{-1}$. The charts (\mathcal{U}, ψ) and (\mathcal{U}', ψ') are said to be *compatible* if the transition map $\psi' \circ \psi^{-1}$ is a homeomorphism (i.e., a continuous bijection with a continuous inverse).

An *atlas* is a collection of compatible charts $\{(\mathcal{U}_\alpha, \psi_\alpha)\}$ which cover the manifold \mathcal{M} . That is, each point $p \in \mathcal{M}$ is in at least one of the chart neighbourhoods \mathcal{U}_α . (Here, α takes values from some indexing set.)

An n -dimensional *topological manifold* consists of a space \mathcal{M} with an atlas $\{(\mathcal{U}_\alpha, \psi_\alpha)\}$ on \mathcal{M} . The manifold is said to be a C^k -*differentiable manifold* if the transition maps $\psi_i \circ \psi_j^{-1}$ are not only continuous but C^k -differentiable. If the transition maps are infinitely differentiable, the manifold is said to be C^∞ (or *smooth*). For differentiable manifolds, the coordinates are related by n differentiable functions with non-vanishing Jacobian at each point of the overlap, i.e., $x^{i'} = x^{i'}(x^j)$ with $\det \left(\frac{\partial x^{i'}}{\partial x^j} \right) \neq 0$ at all points in $\mathcal{U}_i \cap \mathcal{U}_j$.

In addition to the properties listed above, definitions of manifolds often feature additional topological restrictions, such as *Hausdorffness* (i.e., any two distinct points in \mathcal{M} have disjoint neighbourhoods) and *paracompactness* (i.e., every open cover of \mathcal{M} has a locally finite open refinement). A detailed consideration of these properties is not necessary here.

Given two manifolds \mathcal{M} and \mathcal{M}' of dimension n and n' , respectively, the $(n + n')$ -dimensional product space $\mathcal{M} \times \mathcal{M}'$ consisting of all pairs (p, p') with $p \in \mathcal{M}$ and $p' \in \mathcal{M}'$ can be made into a manifold in a natural way. If $\psi_\alpha: \mathcal{U}_\alpha \rightarrow U_\alpha$ and $\psi'_\beta: \mathcal{U}'_\beta \rightarrow U'_\beta$ are charts on \mathcal{M} and \mathcal{M}' , respectively, then one can define a chart on the product $\mathcal{M} \times \mathcal{M}'$ as $\psi_{\alpha\beta}: \mathcal{U}_\alpha \times \mathcal{U}'_\beta \rightarrow U_\alpha \times U'_\beta \subseteq \mathbb{R}^{n+n'}$, by taking $\psi_{\alpha\beta}(p, p') = (\psi_\alpha(p), \psi'_\beta(p'))$. The family of charts $\{\psi_{\alpha\beta}\}$ satisfies the properties required to define a manifold structure on the product $\mathcal{M} \times \mathcal{M}'$.

Having defined a manifold, whose structure is given by charts $\{\psi_\alpha\}$, we may now define the notions of smoothness and differentiability for maps between manifolds. Let \mathcal{M} and \mathcal{M}' be manifolds, and let $\{\psi_\alpha\}$ and $\{\psi'_\beta\}$ denote their respective chart maps. A *map* $f: \mathcal{M} \rightarrow \mathcal{M}'$ is said to be C^∞ (or *smooth*) if, for each pair (α, β) , the map $\psi'_\beta \circ f \circ \psi_\alpha^{-1}$ from \mathbb{R}^n to $\mathbb{R}^{n'}$ is C^∞ .

If a map $f: \mathcal{M} \rightarrow \mathcal{M}'$ is a smooth bijection and has a smooth inverse, then f is called a *diffeomorphism*, and the manifolds \mathcal{M} and \mathcal{M}' are said to be *diffeomorphic*.

A *smooth curve* $\gamma(\lambda)$ in \mathcal{M} is defined to be a smooth map from an interval of \mathbb{R} to \mathcal{M} , where $\lambda \in \mathbb{R}$ is a parameter along the curve.

2.1.2 Vectors, one-forms and tensors

Vectors

Consider an n -dimensional manifold \mathcal{M} . A *tangent vector* v at a point $p \in \mathcal{M}$ is a linear functional $v: C^\infty(\mathcal{M}, \mathbb{R}) \rightarrow \mathbb{R}$, where $C^\infty(\mathcal{M}, \mathbb{R})$ denotes the set of C^∞ functions from \mathcal{M} to \mathbb{R} . The tangent vector is linear and satisfies the Leibniz property:

$$(i) \quad v(c_1 f_1 + c_2 f_2) = c_1 v(f_1) + c_2 v(f_2), \text{ for all } f_1, f_2 \in C^\infty(\mathcal{M}, \mathbb{R}), \text{ and } c_1, c_2 \in \mathbb{R};$$

$$(ii) \quad v(f_1 f_2) = v(f_1) f_2 + f_1 v(f_2).$$

It follows from axioms (i) and (ii) that if $f \in C^\infty(\mathcal{M}, \mathbb{R})$ is a constant function, i.e., $f(p) = \text{constant}$ for all $p \in \mathbb{R}$, then $v(f) = 0$.

A tangent vector is a *directional derivative* along a smooth curve $\gamma(\lambda)$ which passes through p . One may demonstrate by performing a Taylor series expansion of a function f at p and using the axioms (i) and (ii) above that a tangent vector at p may be expressed

in the form

$$v = v^a \frac{\partial}{\partial x^a}. \quad (2.1)$$

The real coefficients v^a are called the *components* of v at the point p , with respect to the local coordinate system $\{x^a\}$ in a neighbourhood of p . The directional derivatives $\frac{\partial}{\partial x^a}$ along coordinate curves at p form a basis of an n -dimensional vector space whose elements are the tangent vectors at p . This vector space is called the *tangent space* (at the point p), and is denoted $T_p\mathcal{M}$. The basis $\left\{\frac{\partial}{\partial x^a}\right\}$ is called a *coordinate basis*; we often write its elements as $\frac{\partial}{\partial x^a} = \partial_a$. In the coordinate basis, the action of a basis vector on a function is written in the form $f_{,a} = \partial_a f = \frac{\partial f}{\partial x^a}$. One may express v in terms of some other basis $\left\{\frac{\partial}{\partial x'^a}\right\}$. The old basis can be expressed in terms of the new basis as

$$\frac{\partial}{\partial x^a} = \frac{\partial x'^b}{\partial x^a} \frac{\partial}{\partial x'^b}. \quad (2.2)$$

Moreover, the components v'^a of v in terms of the new basis are related to those of the old basis by

$$v'^a = \frac{\partial x'^a}{\partial x^b} v^b. \quad (2.3)$$

This is known as the vector transformation law.

One may express a tangent vector in terms of a *general basis* $\{e_a\}$, which is a collection of n linearly independent vectors e_a at p . Any vector $v \in T_p\mathcal{M}$ can then be written in the form

$$v = v^a e_a. \quad (2.4)$$

A coordinate basis is simply a special choice of general basis. Frequently, we will perform calculations in a coordinate basis; however, there will be a number of occasions when it is preferable to use a general basis.

The disjoint union of all tangent spaces $T_p\mathcal{M}$ at points $p \in \mathcal{M}$ forms the *tangent bundle* $T\mathcal{M}$. In local coordinates, the elements of $T\mathcal{M}$ are the $2n$ -tuples (x^a, v^a) . The tangent bundle is then a $2n$ -dimensional manifold. Moreover, if \mathcal{M} is C^k , then $T\mathcal{M}$ is C^{k-1} .

One may now construct a *vector field* $v(p)$ on \mathcal{M} by assigning to each $p \in \mathcal{M}$ a vector $v \in T_p\mathcal{M}$ such that the components v^a are differentiable functions of the local coordinates x^a . A vector field $v(p)$ may then be regarded as a smooth map from \mathcal{M} to $T\mathcal{M}$.

Let $X \in T\mathcal{M}$ be a vector field on \mathcal{M} , and consider a point $p \in \mathcal{M}$. Let $J \subseteq \mathbb{R}$ be an open interval which contains the point 0. A smooth curve $\gamma: J \rightarrow \mathcal{M}$ is called an *integral curve* of X passing through p if it satisfies the initial value problem

$$\dot{\gamma}(\lambda) = X|_{\gamma}, \quad \gamma(0) = p, \quad (2.5)$$

where an overdot denotes differentiation with respect to the parameter λ . In local coordinates $\{x^a\}$, the problem of finding such curves reduces to solving the system

$$\frac{d}{d\lambda} x^a(\lambda) = X^a(x^1(\lambda), \dots, x^n(\lambda)), \quad (2.6)$$

where X^a denotes the a th component of the vector field X in the coordinate basis $\{\partial_a\}$. Given a starting point p at $\lambda = 0$, such a system of ordinary differential equations has a unique solution, so every smooth vector field X has a unique family of integral curves [82].

Let $\mathcal{D} = \{(\lambda, p) \mid \lambda \in J, p \in \mathcal{M}\}$. The *flow of X* is the smooth map $\varphi_X: \mathcal{D} \rightarrow \mathcal{M}$ defined by $\varphi_X(\lambda, p) = \gamma(\lambda)$, where $\gamma: J \rightarrow \mathcal{M}$ is the unique solution to (2.5).

One-forms

A *one-form* (or *dual vector*) σ maps a vector $v \in T_p\mathcal{M}$ into a real number, which is the *contraction* of σ and v , denoted $\langle \sigma, v \rangle$. This mapping is linear, i.e., $\langle \sigma, c_1 v_1 + c_2 v_2 \rangle = c_1 \langle \sigma, v_1 \rangle + c_2 \langle \sigma, v_2 \rangle$, for all $v_1, v_2 \in T_p\mathcal{M}$, and $c_1, c_2 \in \mathbb{R}$. Linear combinations of one-forms are then defined by the linearity property $\langle c_1 \sigma_1 + c_2 \sigma_2, v \rangle = c_1 \langle \sigma_1, v \rangle + c_2 \langle \sigma_2, v \rangle$, for all $c_1, c_2 \in \mathbb{R}$.

The set of n linearly independent one-forms $\{\omega^a\}$, which are determined by $\langle \omega^a, e_b \rangle = \delta^a_b$, form a basis of the *dual space* $T_p^*\mathcal{M}$ to the tangent space $T_p\mathcal{M}$. This is a vector space known as the *cotangent space*. The basis $\{\omega^a\}$ is said to be dual to the basis $\{e_a\}$. A one-form $\sigma \in T_p^*\mathcal{M}$ can be expressed in terms of the basis $\{\omega^a\}$ as

$$\sigma = \sigma_a \omega^a. \quad (2.7)$$

The contraction of any one-form $\sigma \in T_p^*\mathcal{M}$ and any vector $v \in T_p\mathcal{M}$ can be expressed as

$$\langle \sigma, v \rangle = \sigma_a v^a, \quad (2.8)$$

with respect to the bases $\{\omega^a\}$ and $\{e_a\}$.

The *differential* df of a function f is a one-form, with the defining property $\langle df, v \rangle = v(f) = v^a e_a(f)$. Taking f to be the local coordinate functions $\{x^a\}$ and v to be the coordinate basis, the previous definition gives $\langle dx^a, \partial_b \rangle = \frac{\partial x^a}{\partial x^b} = \delta^a_b$, which implies that the basis $\{dx^a\}$ of the cotangent space is dual to the coordinate basis $\{\partial_a\}$ of the tangent space. Of course, any one-form $\sigma \in T_p^*\mathcal{M}$ can then be written in terms of the basis $\{dx^a\}$ as

$$\sigma = \sigma_a dx^a. \quad (2.9)$$

In local coordinates, the differential df is simply $df = f_{,a} dx^a$.

If σ_a denote the components of a one-form with respect to the dual basis $\{dx^a\}$, then it follows from $\langle dx^a, \partial_b \rangle = \delta^a_b$ and the vector transformation law (2.3) that

$$\omega'_a = \frac{\partial x^b}{\partial x'^a} \omega_b. \quad (2.10)$$

This is the transformation law for one-form components.

In a fashion analogous to the construction of the tangent bundle and vector fields, one may construct the *cotangent bundle* $T^*\mathcal{M}$, which is the disjoint union of cotangent spaces $T_p^*\mathcal{M}$. A *one-form field* $\sigma(p)$ is then constructed by assigning to each $p \in \mathcal{M}$ a one-form $\sigma \in T_p^*\mathcal{M}$, such that the components σ_a are differentiable functions of the local coordinates. The components σ_a are often referred to as the covariant components of a vector.

Tensors

If V and W are vector spaces, then the *tensor product* of V and W , denoted $V \otimes W$, is also a vector space. There is a standard bilinear map $V \times W \rightarrow V \otimes W$, denoted by $(v, w) \mapsto v \otimes w$, which satisfies the following axioms:

$$(i) \quad (c_1 v_1 + c_2 v_2) \otimes w = c_1 v_1 \otimes w + c_2 v_2 \otimes w;$$

$$(ii) \quad v \otimes (c_1 w_1 + c_2 w_2) = c_1 v \otimes w_1 + c_2 v \otimes w_2;$$

for all $v_1, v_2 \in V$, $w_1, w_2 \in W$ and $c_1, c_2 \in \mathbb{R}$. The vector space $V \otimes W$ is then the space of all finite linear combinations of formal symbols of the form $v \otimes w$ for $v \in V$ and $w \in W$.

A *tensor* T of type (r, s) at p is a multilinear map from the tensor product of r copies of the tangent space at p with s copies of the cotangent space at p to the real numbers:

$$T: T_p \mathcal{M} \otimes \dots \otimes T_p \mathcal{M} \otimes T_p^* \mathcal{M} \otimes \dots \otimes T_p^* \mathcal{M} \rightarrow \mathbb{R}. \quad (2.11)$$

The tensor T maps any ordered set of r one-forms and s vectors into a real number.

An arbitrary type- (r, s) tensor T can be expressed in terms of the bases $\{e_a\}$ and $\{\omega^b\}$ as

$$T = T^{a_1 \dots a_r}_{b_1 \dots b_s} e_{a_1} \otimes \dots \otimes e_{a_r} \otimes \omega^{b_1} \otimes \dots \otimes \omega^{b_s}. \quad (2.12)$$

The *components* of the tensor T with respect to $\{e_a\}$ and $\{\omega^b\}$ are the real coefficients $T^{a_1 \dots a_r}_{b_1 \dots b_s}$. In general, the components of a tensor T of type (r, s) transform as

$$T'^{a_1 \dots a_r}_{b_1 \dots b_s} = \frac{\partial x'^{a_1}}{\partial x^{c_1}} \dots \frac{\partial x'^{a_r}}{\partial x^{c_r}} \frac{\partial x^{d_1}}{\partial x'^{b_1}} \dots \frac{\partial x^{d_s}}{\partial x'^{b_s}} T^{c_1 \dots c_r}_{d_1 \dots d_s}. \quad (2.13)$$

The generalisation of a tensor to a *tensor field* is straightforward: a tensor field is a choice of a tensor at each point $p \in \mathcal{M}$ that varies smoothly with any coordinates (i.e., the components of the tensor are smooth functions of the local coordinates). In particular, type- $(1, 0)$ tensors are vector fields, type- $(0, 1)$ tensors are one-form (or covector) fields, and type- $(0, 0)$ tensors are defined to be functions. For simplicity, we hereafter refer to tensor fields as tensors.

Exterior calculus

At this stage we introduce some notation for the totally symmetric and totally antisymmetric parts of tensors. For a tensor T of type $(0, p)$ with components $T_{a_1 \dots a_p}$, we define its totally symmetric and totally antisymmetric parts, respectively, as

$$T_{(a_1 \dots a_p)} = \frac{1}{p!} \sum_{\pi} T_{a_{\pi(1)} \dots a_{\pi(p)}}, \quad (2.14)$$

$$T_{[a_1 \dots a_p]} = \frac{1}{p!} \sum_{\pi} \text{sgn}(\pi) T_{a_{\pi(1)} \dots a_{\pi(p)}}, \quad (2.15)$$

where the summation is taken over all permutations π of the ordered set $\{1, \dots, p\}$, and $\text{sgn}(\pi)$ is the sign of the permutation, which takes the value $+1$ (-1) for even (odd) permutations of $\{1, \dots, p\}$. (Equivalent definitions apply for the symmetrisation and anti-symmetrisation of contravariant indices.)

A p -form α is a totally antisymmetric tensor of type $(0, p)$. The set of all p -forms on a manifold \mathcal{M} is a vector space, denoted $\Omega^p(\mathcal{M})$. A smooth function f on \mathcal{M} is a 0-form, $f \in \Omega^0(\mathcal{M})$.

The *exterior product* (or *wedge product*) of a p -form α with a q -form β is denoted by $\alpha \wedge \beta$. It is a totally antisymmetric tensor of type $(0, p + q)$, i.e., a $(p + q)$ -form, whose components are given (up to normalisation) by the anti-symmetrisation of the tensor product of α and β :

$$(\alpha \wedge \beta)_{a_1 \dots a_p b_1 \dots b_q} = \frac{(p+q)!}{p!q!} \alpha_{[a_1 \dots a_p} \beta_{b_1 \dots b_q]}. \quad (2.16)$$

The exterior product obeys the property $\alpha \wedge \beta = (-1)^{pq} \beta \wedge \alpha$.

For one-forms α and β , their *symmetric product* is defined in terms of the tensor product as

$$\alpha\beta = \frac{1}{2}(\alpha \otimes \beta + \beta \otimes \alpha). \quad (2.17)$$

The *exterior derivative* d is a map $d: \Omega^p(\mathcal{M}) \rightarrow \Omega^{p+1}(\mathcal{M})$, which is completely determined by the axioms

- (i) $d(\alpha + \beta) = d\alpha + d\beta$;
- (ii) $d(\alpha \wedge \beta) = d\alpha \wedge \beta + (-1)^p \alpha \wedge d\beta$;
- (iii) $d^2 f = d(df) = 0$;
- (iv) $df = f_{,a} dx^a$.

Property (iv) says that the exterior derivative maps a function f (i.e., a 0-form) to its differential.

A differential p -form α is *closed* if its exterior derivative vanishes ($d\alpha = 0$). A differential p -form α is *exact* if it can be expressed as the exterior derivative of another differential $(p-1)$ -form β ($\alpha = d\beta$). The form β is called a *potential form* for α . The potential form β is non-unique: $\beta' = \beta + d\gamma$, where γ is any $(p-2)$ -form, is also a potential form for α , since $d^2\gamma = 0$.

By property (iii) from the above list of properties for the exterior derivative operator, any exact form is necessarily closed. The question of whether the converse of this statement is true depends on the topology of the domain of interest. On a contractible domain, every closed form is exact by the Poincaré lemma.

2.1.3 Metric tensor

A *pseudo-Riemannian manifold* (\mathcal{M}, g) is an n -dimensional smooth manifold \mathcal{M} endowed with a tensor g of type $(0, 2)$, called the *metric tensor*, such that, at each point $p \in \mathcal{M}$, g is a symmetric non-degenerate bilinear quadratic form. In a coordinate basis, one can write the metric as

$$g = g_{ab} dx^a \otimes dx^b, \quad (2.18)$$

Often, we use the notation ds^2 to denote the *line element*, writing

$$ds^2 = g_{ab} dx^a dx^b, \quad (2.19)$$

where $dx^a dx^b$ denotes the symmetric product of the basis one-forms dx^a and dx^b ; see (2.17). The notation (2.19) is consistent with the intuitive notion that the metric represents an infinitesimal squared distance on the manifold \mathcal{M} . A *Riemannian manifold* is equipped with a metric of signature $(+, +, \dots, +)$; whereas a *Lorentzian manifold* is endowed with a metric of signature $(-, +, \dots, +)$.

Einstein's theory of general relativity is based on the concept of *spacetime*, a four-dimensional continuum which unifies the three spatial dimensions and one temporal dimension. In this work, a *spacetime* is a four-dimensional Lorentzian manifold \mathcal{M} , equipped with a metric g of signature $(-, +, +, +)$.

The scalar product of two vectors v and u is given by the contraction $v \cdot u = g_{ab} v^a u^b$. Two vectors are said to be *orthogonal* if their scalar product vanishes. A vector v is said to be *timelike*, *null* or *spacelike* if $g_{ab} v^a v^b$ is negative, zero or positive, respectively.

The contravariant components g^{ab} form the matrix inverse of the metric tensor g_{ab} ; the "inverse metric" g^{ab} is a type-(2, 0) tensor. One may use the metric and its inverse to raise and lower indices in the standard fashion; for example, $v_a = g_{ab} v^b$ and $v^a = g^{ab} v_b$. This means that the vector field $v^a e_a$ and the one-form field $v_a \omega^a$ represent the same geometrical object. These will be used interchangeably, and we will often denote a vector (one-form) field by its components, i.e., v^a (v_a).

Let $\Omega > 0$ be a smooth function. A *conformal transformation* of the metric is a mapping of the form $g_{ab} \mapsto \tilde{g}_{ab} = \Omega^2 g_{ab}$. The inverse metrics are related by $\tilde{g}^{ab} = \frac{1}{\Omega^2} g^{ab}$. Conformal transformations arise in a range of contexts in general relativity; we will see that they are particularly important in the treatment of geodesics (Section 2.2).

On an n -dimensional pseudo-Riemannian manifold \mathcal{M} , the *Hodge dual* of a p -form α is a $(n - p)$ -form ${}^* \alpha$, whose components are defined by

$$({}^* \alpha)_{a_{p+1} \dots a_n} = \alpha^{a_1 \dots a_p} \epsilon_{a_1 \dots a_p a_{p+1} \dots a_n}, \quad (2.20)$$

where ϵ is the Levi-Civita tensor (or totally anti-symmetric tensor). On four-dimensional spacetime, the Hodge dual of a p -form is a $(4 - p)$ -form, and the Levi-Civita tensor is $\epsilon_{abcd} = \sqrt{-\det g} [abcd]$, where $\det g$ denotes the determinant of the metric tensor, and $[abcd]$ is the fully anti-symmetric Levi-Civita symbol with $[0123] = 1$.

2.1.4 Covariant differentiation and parallel transport

A *covariant derivative* operator ∇ is a map which takes a smooth type- (r, s) tensor field to a smooth type- $(r, s + 1)$ tensor field. For a tensor T with components $T^{a_1 \dots a_r}_{b_1 \dots b_s}$, the action of the covariant derivative on T is denoted $\nabla_c T^{a_1 \dots a_r}_{b_1 \dots b_s}$ in index notation; we employ the standard notation ∇_a , attaching an index to the covariant derivative operator.

On an n -dimensional smooth pseudo-Riemannian manifold (\mathcal{M}, g) , the covariant derivative satisfies the following conditions.

- (i) **Linearity:** $\nabla_c \left(\alpha X^{a_1 \dots a_r}_{b_1 \dots b_s} + \beta Y^{a_1 \dots a_r}_{b_1 \dots b_s} \right) = \alpha \nabla_c X^{a_1 \dots a_r}_{b_1 \dots b_s} + \beta \nabla_c Y^{a_1 \dots a_r}_{b_1 \dots b_s}$, for all tensors X, Y of type (r, s) , and all $\alpha, \beta \in \mathbb{R}$.
- (ii) **Leibniz rule for the tensor product of two tensors:** $\nabla_e \left(X^{a_1 \dots a_r}_{b_1 \dots b_s} Y^{c_1 \dots c_{r'}}_{d_1 \dots d_{s'}} \right) =$

$\left(\nabla_e X^{a_1 \dots a_r}_{b_1 \dots b_s}\right) Y^{c_1 \dots c_{r'}}_{d_1 \dots d_{s'}} + X^{a_1 \dots a_r}_{b_1 \dots b_s} \left(\nabla_e Y^{c_1 \dots c_{r'}}_{d_1 \dots d_{s'}}\right)$, for all tensors X of type (r, s) and tensors Y of type (r', s') .

- (iii) Commutativity with contraction: $\nabla_c \left(T^{a_1 \dots d \dots a_r}_{b_1 \dots d \dots b_s}\right) = \nabla_c T^{a_1 \dots d \dots a_r}_{b_1 \dots d \dots b_s}$, for all tensors T of type (r, s) (so that the parentheses are not necessary).
- (iv) Consistency with tangent vectors as directional derivatives of scalar functions: $v(f) = v^a \nabla_a f$, for all smooth functions $f \in C^\infty(\mathcal{M}, \mathbb{R})$ and all tangent vectors $v \in T\mathcal{M}$.
- (v) Torsion-free: $\nabla_a \nabla_b f = \nabla_b \nabla_a f$, for all smooth functions $f \in C^\infty(\mathcal{M}, \mathbb{R})$.

Imposing the supplementary condition of *metric compatibility* $\nabla_a g_{bc} = 0$, defines a unique covariant derivative operator ∇_a . Using the Levi-Civita connection, we may express the action of the covariant derivative operator ∇_a on a type- (r, s) tensor T in terms of the ordinary partial derivative operator ∂_a in a coordinate basis $\{x^a\}$ as

$$\nabla_a T^{b_1 \dots b_r}_{c_1 \dots c_s} = \partial_a T^{b_1 \dots b_r}_{c_1 \dots c_s} + \sum_i \Gamma^b_{ad} T^{b_1 \dots d \dots b_r}_{c_1 \dots c_s} - \sum_i \Gamma^d_{ac_i} T^{b_1 \dots b_r}_{c_1 \dots d \dots c_s}, \quad (2.21)$$

where there is one term for each contravariant (covariant) index of T which comes with a coefficient of $+1$ (-1).¹ The *connection coefficients* Γ^a_{bc} are known as the *Christoffel symbols*, which are defined by

$$\Gamma^a_{bc} = \frac{1}{2} g^{ad} (\partial_b g_{dc} + \partial_c g_{bd} - \partial_d g_{bc}). \quad (2.22)$$

The Christoffel symbols are symmetric in their lower indices (i.e., $\Gamma^a_{bc} = \Gamma^a_{cb}$), due to the torsion-free property of the covariant derivative. We caution here that the connection coefficients Γ^a_{bc} are *not* tensors as they do not obey the tensor transformation law (2.13) under changes of coordinates. However, the covariant derivative of a tensor field *does* transform covariantly.

Special cases of (2.21) are the action of the covariant derivative on functions f , vectors v^a , and one-forms σ_a ; these are, respectively, given by

$$\nabla_a f = \partial_a f, \quad \nabla_a v^b = \partial_a v^b + \Gamma^b_{ac} v^c, \quad \nabla_a \sigma_b = \partial_a \sigma_b - \Gamma^c_{ab} \sigma_c. \quad (2.23)$$

In general, we will denote the covariant derivative using a semi-colon (and the partial derivative using a comma), e.g. $\nabla_a v^b = v^b_{;a} = v^b_{,a} + \Gamma^b_{ac} v^c$.

Having defined the covariant derivative ∇_a , we may now describe the notion of parallel transport of a vector (or, more generally, a tensor) along a curve γ with tangent vector u^a . A vector v^a at each point on γ is said to be *parallel-transported* (or *parallel-propagated*) along the curve γ if its covariant derivative vanishes along the curve, i.e.,

$$u^b \nabla_b v^a = 0. \quad (2.24)$$

In general, a tensor field of type (r, s) with components $T^{a_1 \dots a_r}_{b_1 \dots b_s}$ is parallel-transported along γ if

$$u^c \nabla_c T^{a_1 \dots a_r}_{b_1 \dots b_s} = 0 \quad (2.25)$$

¹One may, of course, define a covariant derivative using connection other than the Levi-Civita connection; for a more complete discussion, see e.g. Wald [82].

along the curve. In a coordinate basis $\{x^a\}$, one may write the equation of parallel transport of a vector field (2.24) as $u^b \partial_b v^a + \Gamma^a_{bc} u^b v^c = 0$. If the curve is parametrised by a parameter $\lambda \in \mathbb{R}$, then one may use the chain rule to rewrite this as

$$\dot{v}^a + \Gamma^a_{bc} u^b v^c = 0, \quad (2.26)$$

where an overdot denotes differentiation with respect to λ .

2.1.5 Lie differentiation

The *commutator* of two vector fields u and v is itself a vector field, denoted $[u, v]$. It is defined by its action on smooth functions f :

$$[u, v]f = u(v(f)) - v(u(f)). \quad (2.27)$$

For all vector fields u, v and w , the commutator satisfies the anti-symmetry property and the Jacobi identity:

- (i) $[u, v] = -[v, u]$;
- (ii) $[u, [v, w]] + [v, [w, u]] + [w, [u, v]] = 0$.

On an n -dimensional manifold \mathcal{M} , one may use the commutator of vector fields to define the *Lie derivative* of a type- (r, s) tensor field T along the vector field v , which is itself a type- (r, s) tensor field, denoted $\mathcal{L}_v T$. The Lie derivative evaluates the change of a tensor field along the flow induced by a vector field.

For smooth scalar functions f , the Lie derivative of f along the vector field u is given by

$$\mathcal{L}_u f = u(f) = u^a \partial_a f. \quad (2.28)$$

The Lie derivative of a vector field v along the vector field u is defined by

$$\mathcal{L}_u v = [u, v]. \quad (2.29)$$

This definition may be extended to tensors of arbitrary type, by demanding that the Lie derivative satisfies the following properties.

- (i) **Linearity:** $\mathcal{L}_u(\alpha X + \beta Y) = \alpha \mathcal{L}_u X + \beta \mathcal{L}_u Y$, for all type- (r, s) tensors X, Y , and all $\alpha, \beta \in \mathbb{R}$.
- (ii) **Leibniz rule on the tensor product of two tensors:** $\mathcal{L}_u(X \otimes Y) = (\mathcal{L}_u X) \otimes Y + X \otimes (\mathcal{L}_u Y)$ for all tensors X, Y .
- (iii) **“Product rule” on the contraction of a vector and a one-form:** $\mathcal{L}_u \langle \sigma, v \rangle = \langle \mathcal{L}_u \sigma, v \rangle + \langle \sigma, \mathcal{L}_u v \rangle$, for all one-forms σ and vectors v .

The above properties may be used to deduce the components of the Lie derivative of an arbitrary type- (r, s) tensor field:

$$(\mathcal{L}_u T)^{a_1 \dots a_r}_{b_1 \dots b_s} = u^c T^{a_1 \dots a_r}_{b_1 \dots b_s, c} - \sum_i T^{a_1 \dots c \dots a_r}_{b_1 \dots b_s} u^{a_i}, c + \sum_i T^{a_1 \dots a_r}_{b_1 \dots c \dots b_s} u^c, b_i, \quad (2.30)$$

where there is one term for each index, and the covariant (contravariant) indices come with a coefficient of +1 (−1).

For the metric tensor g_{ab} , one may replace the partial derivatives on the right-hand side of (2.30) with covariant derivatives using the Levi-Civita connection. This gives the Lie derivative of the metric tensor along the vector field u :

$$(\mathcal{L}_u g)_{ab} = u^c g_{ab;c} + g_{cb} u^c{}_{;a} + g_{ac} u^c{}_{;b} = u_{a;b} + u_{b;a}, \quad (2.31)$$

where we have used the metric compatibility of the covariant derivative, i.e., $g_{ab;c} = 0$.

A tensor field T of type (r, s) is said to be *Lie-transported* along the curve γ with tangent vector field u if its Lie derivative along the curve vanishes, i.e., $\mathcal{L}_u T = 0$.

We shall see later that the Lie derivative plays an important role in describing the symmetries of the gravitational field in general relativity.

2.1.6 Killing vectors

Symmetries are ubiquitous in physics and play a fundamental role in general relativity. Here, we briefly describe explicit continuous symmetries of spacetime. A spacetime with metric g admits a *continuous symmetry (isometry)* if there exists a continuous transformation of the spacetime into itself which preserves the metric.

Such transformations are encoded by *Killing vectors* ξ ; the isometry condition states that the Lie derivative of the metric with respect to ξ vanishes, i.e.,

$$\mathcal{L}_\xi g = 0. \quad (2.32)$$

Geometrically, the condition $\mathcal{L}_\xi g = 0$ says that the metric is invariant under the flow of ξ .

Using the metric connection ∇_a in local coordinates, we see from (2.31) that the isometry condition is equivalent to the *Killing vector equation*,

$$\nabla_a \xi_b + \nabla_b \xi_a = 0. \quad (2.33)$$

This equation is often written in the form $\xi_{(a;b)} = 0$.

If the metric coefficients g_{ab} in some basis $\{dx^a\}$ are independent of one of the coordinate x^α (for some fixed α), then $\xi^\alpha = \delta^\alpha_\alpha$ is a Killing vector field in the basis $\{\partial_a\}$. To see this, we first note that $g_{ab,\alpha} = 0$. We have $\xi^\alpha = \delta^\alpha_\alpha$, so $\xi_a = g_{ab} \xi^b = g_{ab} \delta^b_\alpha = g_{a\alpha}$. The left-hand side of the Killing vector equation (2.33) is then

$$\xi_{a;b} + \xi_{b;a} = \xi_{a,b} + \xi_{b,a} - 2\Gamma^c{}_{ab} \xi_c = g_{a\alpha,b} + g_{b\alpha,a} - g^{cd} (g_{bd,a} + g_{ad,a} - g_{ab,d}) g_{c\alpha}. \quad (2.34)$$

Now, using the fact that $g^{cd} g_{c\alpha} = \delta^d_\alpha$, the right-hand side of (2.34) becomes

$$g_{a\alpha,b} + g_{b\alpha,a} - (g_{b\alpha,a} + g_{a\alpha,a} - g_{ab,\alpha}) = -g_{ab,\alpha} = 0, \quad (2.35)$$

by virtue of the fact that g_{ab} is independent of x^α . Hence, ξ^α is a Killing vector.

2.1.7 Stationary and static spacetimes

Consider a spacetime (\mathcal{M}, g_{ab}) with local coordinates x^a . A non-vanishing one-form field u_a is said to be *hypersurface-orthogonal* if it is orthogonal to a family of $\Phi = \text{constant}$ hypersurfaces for some function Φ with $\partial_a \Phi \neq 0$. In other words, $u_a = \alpha \partial_a \Phi$ for some non-zero function of the spacetime coordinates α . The Frobenius theorem says that, locally, the definition of hypersurface-orthogonality is equivalent to $u_{[a} \partial_b u_{c]} = u_{[a} \nabla_b u_{c]} = 0$. A vector field u^a is hypersurface-orthogonal if and only if the corresponding one-form field $u_a = g_{ab} u^b$ is hypersurface-orthogonal.

A spacetime is *stationary* if it admits a timelike Killing vector ξ . If this is the case, there exist local coordinates (t, x^i) in which $\xi = \partial_t$, $g_{tt} < 0$, and $g_{ab,t} = 0$ (i.e., the metric is independent of t). The off-diagonal terms g_{ti} need not vanish.

A spacetime is *static* if (i) it admits a timelike hypersurface-orthogonal Killing vector ξ , and (ii) it can be globally foliated by hypersurfaces that are orthogonal to ξ . The first of these criteria ensures that such hypersurfaces exist locally, but not necessarily globally. In this case, there exist local coordinates (t, x^i) in which $\xi = \partial_t$, $g_{tt} < 0$, $g_{ab,t} = 0$, and $g_{ti} = 0$.

2.1.8 Curvature

The *Riemann curvature tensor*, which has components $R^a{}_{bcd}$, is a type-(1, 3) tensor which maps a one-form σ and three vectors w , u and v into a real number:

$$R^a{}_{bcd} \sigma_a w^b u^c v^d = \langle \sigma, (\nabla_u \nabla_v - \nabla_v \nabla_u - \nabla_{[u,v]}) w \rangle \quad (2.36)$$

$$= \sigma_a \left[u^c (v^d w^a{}_{;d})_{;c} - v^c (u^d w^a{}_{;d})_{;c} - (u^c v^d{}_{;c} - v^c u^d{}_{;c}) w^a{}_{;d} \right] \quad (2.37)$$

$$= \sigma_a (w^a{}_{;dc} - w^a{}_{;cd}) u^c v^d. \quad (2.38)$$

Since σ , u and v are arbitrary, we arrive at the condition

$$R^a{}_{bcd} w^b = w^a{}_{;dc} - w^a{}_{;cd}, \quad (2.39)$$

which is known as the *Ricci identity*.

One may use the expression for the covariant derivative of a tensor (2.21) to write the components of the Riemann tensor in terms of a coordinate basis as

$$R^a{}_{bcd} = \partial_c \Gamma^a{}_{bd} - \partial_d \Gamma^a{}_{bc} + \Gamma^e{}_{bd} \Gamma^a{}_{ec} - \Gamma^e{}_{bc} \Gamma^a{}_{ed}. \quad (2.40)$$

The components $R_{abcd} = g_{ae} R^e{}_{bcd}$ obey the symmetry relations

$$R_{abcd} = -R_{abdc} = -R_{bacd}, \quad R_{abcd} = R_{cdab}, \quad R_{a[bcd]} = 0. \quad (2.41)$$

The last of these is known as the *first (algebraic) Bianchi identity*. The components of the covariant derivative of the Riemann tensor obey the *second (differential) Bianchi identity*,

$$R_{ab[cd;e]} = 0. \quad (2.42)$$

The *Ricci curvature tensor* R_{ab} is defined as the trace of the Riemann tensor over its first and third indices:

$$R_{ab} = R^c{}_{acb}. \quad (2.43)$$

The Ricci tensor satisfies the symmetry property $R_{ab} = R_{ba}$. The *Ricci scalar* R is defined as the trace of the Ricci tensor:

$$R = R^a{}_a. \quad (2.44)$$

The trace-free part of the Riemann tensor is the *Weyl tensor*. For n -dimensional manifolds with $n \geq 3$, the Weyl tensor has components

$$C_{abcd} = R_{abcd} - \frac{2}{n-2}(g_{a[c}R_{d]b} - g_{b[c}R_{d]a}) + \frac{2}{(n-1)(n-2)}Rg_{a[c}g_{b]d}. \quad (2.45)$$

The Weyl tensor satisfies the same symmetry properties as the Riemann tensor, and is trace-free on all of its indices.

The *Einstein tensor* is defined as

$$G_{ab} = R_{ab} - \frac{1}{2}Rg_{ab}. \quad (2.46)$$

Clearly, the Einstein tensor is symmetric, i.e., $G_{ab} = G_{ba}$. Taking the trace of both sides of (2.46), we see that $G^a{}_a = \frac{2-n}{2}R$. In four dimensions, this is $G^a{}_a = -R$. The Einstein tensor is therefore referred to as the *trace-reversed Ricci tensor*. By contracting the second Bianchi identity (2.42), one may show that the Einstein tensor is divergence-free, i.e.,

$$\nabla^a G_{ab} = 0. \quad (2.47)$$

2.2 Geodesics

We now turn to a topic of central importance to this thesis – the notion of geodesic motion. Intuitively, a geodesic is a “straight line” on a curved manifold. Here, we present a formal definition of a geodesic, derive the geodesic equation, discuss causal geodesics, and present the Lagrangian and Hamiltonian formulations of geodesic motion on curved manifolds. We restrict our attention to the special case of four-dimensional spacetime manifolds \mathcal{M} , endowed with a metric g_{ab} of Lorentzian signature, although the concepts reviewed here can be generalised to arbitrary n -dimensional pseudo-Riemannian manifolds.

2.2.1 Geodesic equation

Geodesic curves, parallel transport and affine parameters

On spacetime, a *geodesic* is a curve γ whose tangent vector u^a is parallel-transported along itself, i.e.,

$$u^b \nabla_b u^a = 0. \quad (2.48)$$

In fact, for a curve on a manifold to be a geodesic, we only require that the parallel-transport of the tangent vector points in the same direction as the tangent vector at each point along the curve; the tangent vector does not necessarily have to be of the same length. This yields the weaker condition

$$u^b \nabla_b u^a = \alpha u^a, \quad (2.49)$$

where α is some arbitrary (smooth) function along the curve. One may show that, given a curve which satisfies (2.49), it is always possible to reparametrise so that it satisfies the

stronger condition (2.48). The latter is referred to as an *affine parametrisation*, and can be used without loss of generality. Throughout this work, we shall define a geodesic to be a curve which satisfies (2.48).

In some coordinate system, an affinely parametrised geodesic $\gamma(\lambda)$ is a curve $x^a(\lambda)$. Using (2.26), the tangent vector $u^a = \dot{x}^a$ to the geodesic satisfies

$$\dot{u}^a + \Gamma^a_{bc} u^b u^c = 0, \quad (2.50)$$

where an overdot denotes differentiation with respect to the affine parameter λ . The components of the tangent vector are related to the coordinates by $u^a = \dot{x}^a$, so the geodesic equation can be written

$$\ddot{x}^a + \Gamma^a_{bc} \dot{x}^b \dot{x}^c = 0. \quad (2.51)$$

This is actually a coupled system of four second-order ordinary differential equations for each of the spacetime coordinates $x^a(\lambda)$. For every initial value $x^a(0)$ and $\dot{x}^a(0)$, a unique solution to this system is guaranteed to exist. In other words, given a point $p \in \mathcal{M}$ and a tangent vector $u^a \in T_p\mathcal{M}$, there always exists a unique geodesic through p with tangent vector u^a .

Timelike, null and spacelike geodesics

A geodesic γ with tangent u^a is said to be *timelike* if $g_{ab}u^a u^b < 0$ everywhere along γ ; it is said to be *null* if $g_{ab}u^a u^b = 0$ everywhere along γ ; or it is said to be *spacelike* if $g_{ab}u^a u^b > 0$ everywhere along γ . The first two types of geodesic are referred to as *causal* geodesics. For timelike geodesics, the *proper time* τ measured along the curve is defined to be

$$\tau = \int_{\gamma} \sqrt{-g_{ab}u^a u^b} d\lambda. \quad (2.52)$$

Similarly, along a spacelike geodesic, the *proper length* ℓ is defined as

$$\ell = \int_{\gamma} \sqrt{g_{ab}u^a u^b} d\lambda. \quad (2.53)$$

It is straightforward to show that the proper time and proper length are independent of the choice of parametrisation. For null geodesics, there is no invariant notion of proper time or proper length: a null geodesic can be said to have zero proper time and proper length.

Along an affinely parametrised geodesic with tangent vector $u^a(\lambda)$, the norm $g_{ab}u^a u^b = u^a u_a$ is constant:

$$\frac{d}{d\lambda} (u^a u_a) = u^b (u^a u_a)_{;b} = u^b u^a_{;b} u_a + u^a u^b u_{a;b} = 0. \quad (2.54)$$

For null geodesics, the norm is always zero. For timelike (spacelike) geodesics which are parametrised by proper time (proper length), the norm is -1 ($+1$).

In general relativity, a freely falling *massive* particle follows a timelike geodesic, whereas a freely propagating *massless* particle (e.g. a photon) follows a null geodesic. The terms null geodesic, lightlike geodesic and light ray are used interchangeably.

Geodesic completeness

A geodesic $x^a(\lambda)$ for which the affine parameter λ ranges over all of \mathbb{R} is said to be *complete*. A geodesic that is not complete is called *incomplete*. A spacetime (\mathcal{M}, g) is said to be *geodesically complete* if every geodesic can be extended into a complete geodesic. A spacetime (\mathcal{M}, g) is said to be *timelike* (respectively *null* or *spacelike*) *geodesically complete* if every *timelike* (respectively *null* or *spacelike*) geodesic can be extended into a complete one.

The notion of geodesic (in)completeness is central to the definition of singularities in general relativity. If a spacetime (\mathcal{M}, g) is geodesically incomplete, and it is not possible to extend (\mathcal{M}, g) into a larger spacetime that is geodesically complete, we say that (\mathcal{M}, g) is *singular*. If any scalar quantity which is polynomial in $R^a{}_{bcd}$ or its covariant derivatives diverges along an incomplete geodesic, the spacetime is singular. In this case, the spacetime is said to have a *scalar curvature singularity*. Alternatively, a component of $R^a{}_{bcd}$ or its covariant derivative(s) in a parallel-transported tetrad (see Section 2.5.1) may blow up; this is referred to as a parallel-propagated curvature singularity [82].

Conformal transformations

Consider a conformal transformation of the metric $g_{ab} \mapsto \tilde{g}_{ab} = \Omega^2 g_{ab}$, where $\Omega > 0$ is a smooth function of the spacetime coordinates. Clearly, a conformal transformation preserves causal structure; that is, if a vector is a timelike, null or spacelike with respect to the metric g_{ab} , then it has the same property with respect to the conformally related metric \tilde{g}_{ab} .

Let us consider the geodesics of the conformally related spacetimes (\mathcal{M}, g_{ab}) and $(\mathcal{M}, \tilde{g}_{ab})$. Using the properties of the covariant derivative (Section 2.1.4), one can show that, for an arbitrary vector field v^a , the covariant derivatives ∇_a and $\tilde{\nabla}_a$ are related by [82]

$$\tilde{\nabla}_a v^b = \nabla_a v^b + D^b{}_{ac} v^c, \quad (2.55)$$

where the *connection coefficients* are given by

$$D^a{}_{bc} = \frac{1}{2} \tilde{g}^{ad} (\nabla_b \tilde{g}_{dc} + \nabla_c \tilde{g}_{bd} - \nabla_d \tilde{g}_{bc}). \quad (2.56)$$

This is a generalisation of the relationship between the covariant derivative operator and the partial derivative operator, where the connection coefficients are simply the Christoffel symbols (2.22). Using the Leibniz rule and metric compatibility of ∇_a , we find

$$\nabla_a \tilde{g}_{bc} = \nabla_a (\Omega^2 g_{bc}) = 2\Omega g_{bc} \nabla_a \Omega + \Omega^2 \nabla_a g_{bc} = 2\Omega g_{bc} \nabla_a \Omega. \quad (2.57)$$

The connection coefficients (2.56) can therefore be expressed in terms of g_{ab} and Ω as

$$D^a{}_{bc} = 2\delta^a{}_{(b} \nabla_{c)} \ln \Omega - g_{bc} g^{ad} \nabla_d \ln \Omega. \quad (2.58)$$

Let $\gamma(\lambda)$ be a geodesic of the spacetime (\mathcal{M}, g_{ab}) with tangent vector u^a . Using (2.55) and (2.58), we obtain

$$u^b \tilde{\nabla}_b u^a = u^b \nabla_b u^a + D^a{}_{bc} u^b u^c = 2 \left(u^b \nabla_b \ln \Omega \right) u^a - \left(g_{bc} u^b u^c \right) g^{ad} \nabla_d \ln \Omega. \quad (2.59)$$

In general, we see that γ is *not* a geodesic with respect to the metric \tilde{g}_{ab} ; however, if γ is a *null* geodesic, then the term $g_{bc}u^b u^c$ on the right-hand side of (2.59) will vanish. This leaves us with $u^b \tilde{\nabla}_b u^a = 2(u^b \nabla_b \ln \Omega) u^a$, which is nothing more than the equation for a non-affinely-parametrised geodesic (2.49). This shows that null geodesics are conformally invariant.

2.2.2 Lagrangian formulation

One may derive the geodesic equation (2.51) using the calculus of variations and Hamilton's principle of least action [84]. Our starting point is the *action functional*, which is taken to be the spacetime distance between two fixed endpoints,

$$S[x^a(\lambda)] = \int L' d\lambda, \quad L' = \sqrt{|g_{ab}\dot{x}^a \dot{x}^b|}, \quad (2.60)$$

where L' is the *Lagrangian function*, and an overdot denotes differentiation with respect to the parameter λ . This definition is somewhat degenerate, as the action is invariant under arbitrary reparametrisations of the curve. The requirement that the action functional (2.60) is stationary (i.e., $\delta S = 0$) yields the Euler–Lagrange equations of motion

$$\frac{d}{d\lambda} \left(\frac{\partial L'}{\partial \dot{x}^a} \right) - \frac{\partial L'}{\partial x^a} = 0. \quad (2.61)$$

Inserting the Lagrangian $L' = \sqrt{|g_{ab}\dot{x}^a \dot{x}^b|}$ into the Euler–Lagrange equations gives the equation for non-affinely parametrised geodesics (2.49). Demanding that the curve be affinely parametrised, we arrive at the geodesic equation (2.51). Geodesics are thus space-time paths $x^a(\lambda)$ which extremise the action (2.60).

The Lagrangian function (2.60) is not unique. In fact, any function which yields the same Euler–Lagrange equations of motion is valid. For example, the geodesic equation (2.51) also follows from extremising the action

$$S[x^a(\lambda)] = \int L d\lambda, \quad L = \frac{1}{2} g_{ab} \dot{x}^a \dot{x}^b. \quad (2.62)$$

Unlike (2.60), the action (2.62) is extremal *only* for affinely parametrised curves. We favour the action (2.62) for two reasons: (i) there is no degeneracy associated with reparametrisations of the curve; (ii) the Lagrangian is more tractable, as it does not feature a square root. We note here that the Lagrangian itself is constant along geodesics, with $L = 0$ in the null case, which will be of particular interest in this work.

The flow determined by the Euler–Lagrange equations (2.61) with Lagrangian (2.60) or (2.62) is called the *geodesic flow*. The geodesics are then the projections of the integral curves of the geodesic flow onto the manifold \mathcal{M} . Physically speaking, the geodesic flow governs the motion of a particle that is not subject to any external forces, therefore moving freely on the manifold \mathcal{M} .

In the Lagrangian picture, the geodesic flow is a dynamical system on the tangent bundle $T\mathcal{M}$; for each value of the parameter λ , the system is completely determined in local coordinates by the position vector $x^a(\lambda)$ and the tangent vector (or *velocity*) $\dot{x}^a(\lambda)$. The cotangent bundle is interpreted as the *configuration space* of Lagrangian mechanics.

2.2.3 Hamiltonian formulation

With a change of perspective, the geodesic flow can be viewed as a dynamical system on the cotangent bundle $T^*\mathcal{M}$, where the state of the system is determined in local coordinates by the position vector $x^a(\lambda)$ and the momentum one-form $p_a(\lambda)$; this is sometimes referred to as the *cogeodesic flow*. In this picture, the geodesic flow becomes a Hamiltonian flow on *phase space*, where the latter is the cotangent bundle. One may pass from the Lagrangian formulation to the Hamiltonian formulation using the following standard procedure.

First, one may write down the *Hamiltonian function* H (which is the Legendre transform of the Lagrangian function L), and the *conjugate momenta* p_a ; these take the form

$$H = \dot{x}^a p_a - L, \quad p_a = \frac{\partial L}{\partial \dot{x}^a}, \quad (2.63)$$

where it is understood that the velocities \dot{x}^a are to be replaced by the momenta p_a in the definition of the Hamiltonian H . Using the Lagrangian (2.62), the Hamiltonian for geodesics is given by

$$H = \frac{1}{2} g^{ab} p_a p_b, \quad (2.64)$$

where the conjugate momenta are

$$p_a = g_{ab} \dot{x}^a. \quad (2.65)$$

Clearly the momenta p_a are the covariant components of the tangent vector to the geodesic. Moreover, the Hamiltonian (2.64) is manifestly covariant.

A standard variational approach yields *Hamilton's equations*, a system of $2n$ coupled first-order ordinary differential equations for each of the n spacetime coordinates x^a and their conjugate momenta p_a , given by

$$\dot{x}^a = \frac{\partial H}{\partial p_a}, \quad \dot{p}_a = -\frac{\partial H}{\partial x^a}. \quad (2.66)$$

Using the Hamiltonian (2.64), Hamilton's equations (2.66) become

$$\dot{x}^a = g^{ab} p_b, \quad \dot{p}_a = -\frac{1}{2} g^{bc}{}_{,a} p_b p_c. \quad (2.67)$$

The first of these is clearly equivalent to the definition of the conjugate momenta (2.65). The equations (2.67) may be combined to give the familiar geodesic equation $p^b \nabla_b p_a = 0$.

We remark here that the Hamiltonian function (2.64) is conserved along geodesics. In the case of timelike geodesics, i.e., the paths of freely falling massive particles, we have $H = -\frac{1}{2}m^2$, where m is the mass of the particle. In the null case, we see that the Hamiltonian vanishes along null rays. In general, the constancy of the Hamiltonian is referred to as the *Hamiltonian constraint*; for null geodesics, the constraint $H = 0$ is referred to as the *null condition*.

In Section 2.1.6, we saw that continuous spacetime symmetries are encoded by Killing vectors. These spacetime symmetries can always be “lifted up” to symmetries on phase space. In Section 2.6.4, we will discuss the relationship between Killing vectors (or, more generally, higher-rank Killing tensors) and conserved quantities on phase space.

2.2.4 Geodesic deviation equation

Our final task for this section is to consider the behaviour of a *congruence* (or *family*) of non-intersecting geodesics. Let $x^a(\lambda, \sigma)$ be a congruence of geodesics, where λ is an affine parameter along each curve, and σ parametrises the congruence, i.e., each geodesic is distinguished by the value of the parameter σ . The tangent vector field is given by

$$u^a = \frac{\partial x^a}{\partial \lambda}, \quad (2.68)$$

which satisfies the geodesic equation $u^b \nabla_b u^a = 0$. We define the *separation vector* (or *deviation vector*) between neighbouring rays in the congruence to be

$$\xi^a = \frac{\partial x^a}{\partial \sigma}. \quad (2.69)$$

It is straightforward to show, using the symmetry of mixed partial derivatives, that ξ is Lie-transported along each geodesic in the congruence:

$$\mathcal{L}_u \xi^a = [u, \xi]^a = u^b \partial_b \xi^a - \xi^b \partial_b u^a = \frac{\partial^2 x^a}{\partial \lambda \partial \sigma} - \frac{\partial^2 x^a}{\partial \sigma \partial \lambda} = 0. \quad (2.70)$$

It follows from the torsion-free property of the Christoffel symbols ($\Gamma^a_{[bc]} = 0$) that the partial derivatives in (2.70) may be replaced by covariant derivatives. We then find

$$u^b \nabla_b \xi^a = \xi^b \nabla_b u^a, \quad (2.71)$$

which says that the parallel transport of ξ along u is equivalent to the parallel transport of u along ξ . A consequence of this result is that the quantity $\xi^a u_a$ is conserved along each geodesic, i.e., $\frac{d}{d\lambda}(\xi^a u_a) = 0$.

Recall that the commutator of covariant derivatives acting on a vector field satisfies

$$\nabla_a \nabla_b u^c - \nabla_b \nabla_a u^c = -R^c_{\quad dab} u^d, \quad (2.72)$$

where $R^a_{\quad bcd}$ are the components of the Riemann curvature tensor. Applying the directional derivative operator $u^c \nabla_c$ to both sides of (2.71) and using the above results, we see that the acceleration of the deviation vector ξ^a is given by

$$\frac{D^2 \xi^a}{d\lambda^2} = R^a_{\quad bcd} u^b u^c \xi^d, \quad (2.73)$$

where $\frac{D}{d\lambda} = u^a \nabla_a$ denotes the covariant derivative along u^a . This equation is known as the *geodesic deviation equation* (or *Jacobi equation*). It describes how spacetime curvature (i.e., the Riemann curvature term on the right-hand side) is responsible for the relative acceleration of neighbouring geodesics in a congruence, even if the rays are initially parallel.

2.3 Field equations and solutions

2.3.1 Einstein field equations

The cornerstone of Einstein's theory of general relativity is Einstein's field equation, which couples the (dynamical) spacetime background to the matter fields present in spacetime.

The field equations may be derived using a principle of least action. Our starting point is the action functional

$$S = \int (\mathcal{L}_{\text{EH}} + \mathcal{L}_{\text{M}}) \sqrt{-\det g} \, d^4x, \quad \mathcal{L}_{\text{EH}} = \frac{1}{16\pi} R, \quad (2.74)$$

where \mathcal{L}_{EH} is the *Einstein–Hilbert Lagrangian* for the gravitational field; \mathcal{L}_{M} is the *matter Lagrangian*, which contains information about the matter fields present in spacetime; $\det g$ is the metric determinant; and R is the Ricci curvature scalar. Variation of the action (2.74) yields the *Einstein field equation*,

$$G_{ab} = 8\pi T_{ab}, \quad (2.75)$$

where $G_{ab} = R_{ab} - \frac{1}{2}Rg_{ab}$ is the Einstein tensor (see Section 2.1.8), and T_{ab} is the *stress–energy tensor*. Here, we outline the key properties of the field equation (2.75). Firstly, the field equation is tensorial; in fact, there are ten independent components, due to the fact that the Einstein tensor and the stress–energy tensor are symmetric. We will often refer to (2.75) as the Einstein field equations (using the plural) to reflect this. Secondly, the field equations constitute a system of coupled second-order differential equations for the components of the metric tensor g_{ab} . Thirdly, the fact that the Einstein tensor is divergence-free (see Section 2.1.8) means that the stress–energy tensor also satisfies the same property, i.e., $\nabla^a T_{ab} = 0$. Finally, the constant of proportionality between the Einstein tensor and stress–energy tensor in (2.75) – which in non-geometrised units is $\frac{8\pi G}{c^2}$, where G is Newton’s gravitational constant and c is the speed of light *in vacuo* – is chosen so that we recover Newton’s theory of gravitation in the limit of slowly varying and weak gravitational fields.

Taking the trace of (2.75) with respect to the metric gives $-R = 8\pi T$, where $T = T^a_a$ denotes the trace of the stress–energy tensor. Replacing the term involving R in (2.75), one obtains the *trace-reversed* form of the Einstein field equations,

$$R_{ab} = 8\pi \left(T_{ab} - \frac{1}{2} T g_{ab} \right). \quad (2.76)$$

In pure vacuum, the stress–energy tensor vanishes, i.e., $T_{ab} = 0$. It is straightforward to see from (2.75) that the vacuum field equations may be expressed in the form $G_{ab} = 0$. Equivalently, by setting $T_{ab} = 0$ in the trace-reversed field equations (2.76), the vacuum field equations can be written in the form $R_{ab} = 0$. Solutions to the vacuum field equations of general relativity are referred to as *vacuum solutions*. The simplest example of such a solution is Minkowski spacetime (see Section 2.4.1). The Schwarzschild and Kerr black hole spacetimes are also examples of vacuum solutions (see Section 2.4.3). In general, manifolds with a vanishing Ricci tensor ($R_{ab} = 0$) are referred to as *Ricci-flat* manifolds.

2.3.2 Einstein–Maxwell field equations

The Einstein–Maxwell field equations of gravitation and electromagnetism may be obtained with the inclusion of the free-field Lagrangian $\mathcal{L}_{\text{EM}} = -\frac{1}{4}F_{ab}F^{ab}$ in the action (2.74), where F_{ab} is the *Faraday tensor*, an antisymmetric rank-two tensor (i.e., a two-form) which encodes the electromagnetic field. The Euler–Lagrange equations for the new

action yield the *Einstein–Maxwell equations*

$$G_{ab} = 8\pi T_{ab}, \quad T_{ab} = \frac{1}{4\pi} \left(F_{ac} F_b{}^c - \frac{1}{4} g_{ab} F^{cd} F_{cd} \right), \quad (2.77)$$

$$\nabla_b F^{ab} = \mu_0 J^a, \quad \nabla_b {}^*F^{ab} = 0, \quad (2.78)$$

where ${}^*F^{ab}$ is the Hodge dual of the Faraday tensor; and J^a is the four-current, which is divergence-free ($\nabla_a J^a = 0$). In the absence of sources ($J^a = 0$), the first of Maxwell's equations (2.78) is simply $\nabla_b F^{ab} = 0$.

In the language of differential forms, the Faraday two-form is *closed*, i.e., $dF = 0$, so F must be locally exact ($F = dA$ for some one-form potential A) by Poincaré's lemma. Hence, the components of the Faraday tensor can be expressed in terms of the one-form potential A_a as

$$F_{ab} = \nabla_a A_b - \nabla_b A_a. \quad (2.79)$$

Taking the trace of the electromagnetic stress–energy tensor (2.77), it is clear that $T = T^a{}_a = 0$. Using the trace-reversed field equations (2.76), we see that

$$R_{ab} = 8\pi T_{ab} = 2 \left(F_{ac} F_b{}^c - \frac{1}{4} g_{ab} F^{cd} F_{cd} \right) \quad (2.80)$$

for pure electromagnetism.

2.3.3 Stationary axisymmetric solutions

A solution to the Einstein–Maxwell equations which will be of central importance to this thesis is that of a rotating source with a stationary, axially symmetric gravitational field. The most general such solution is the *Weyl–Lewis–Papapetrou geometry*, which is described in Weyl–Lewis–Papapetrou coordinates $\{t, \rho, z, \phi\}$ by the line element

$$ds^2 = g_{ab} dx^a dx^b = -f (dt + w d\phi)^2 + \frac{1}{f} [e^{2\gamma} (d\rho^2 + dz^2) + \rho^2 d\phi^2], \quad (2.81)$$

and electromagnetic one-form potential

$$A_a dx^a = A_t dt + A_\phi d\phi, \quad (2.82)$$

where f , γ , w , A_t and A_ϕ are functions of the coordinates ρ and z only.

The metric represents a stationary, axisymmetric solution as the metric (2.81) admits a timelike Killing vector ∂_t and a spacelike Killing vector ∂_ϕ , and the coordinate ϕ is taken to be periodic with $\phi \in [0, 2\pi)$. The function w which appears in (2.81) represents a rotation about the symmetry axis $\rho = 0$. In the case $w = 0$, we recover the so-called *Weyl solution* for a *static* axisymmetric gravitational field. In pure vacuum, the one-form potential (2.82) vanishes.

For the Weyl–Lewis–Papapetrou metric (2.81) with one-form potential (2.82), the electrovacuum field equations are non-linear, but are known to be integrable nevertheless. One may use the Ernst formulation in terms of complex potentials to solve the field equations [83, 85, 86].

2.4 Black holes

2.4.1 Minkowski spacetime and asymptotic flatness

The simplest solution to Einstein's vacuum field equations is *Minkowski spacetime* (or *flat spacetime*), which is described in Cartesian coordinates $\{t, x, y, z\}$ by the line element

$$ds^2 = \eta_{ab} dx^a dx^b = -dt^2 + dx^2 + dy^2 + dz^2, \quad (2.83)$$

where we use the standard notation η_{ab} to denote the components of the Minkowski metric.

We may use the Minkowski metric to provide a coordinate-dependent definition of asymptotic flatness. A spacetime (\mathcal{M}, g_{ab}) is *asymptotically flat* if there exist some coordinates $\{t, x, y, z\}$ such that the metric can be expressed in the form $g_{ab} = \eta_{ab} + h_{ab}$, with $\lim_{r \rightarrow \infty} h_{ab} = O(\frac{1}{r})$, where $r^2 = x^2 + y^2 + z^2$. For a more rigorous, coordinate-invariant definition of asymptotic flatness, we refer the reader to Wald [82].

2.4.2 Event horizon

Consider a spacetime (\mathcal{M}, g_{ab}) . The infinities of this spacetime are identified as follows. *Future (past) null infinity*, denoted \mathcal{I}^+ (\mathcal{I}^-), is the future (past) infinity of all null curves. *Future (past) timelike infinity*, denoted i^+ (i^-), is the future (past) infinity of all timelike curves. *Spacelike infinity*, denoted i^0 , is the future (past) infinity of all spacelike curves. For a more detailed description of these infinities, and the notion of a *conformal boundary*, see [82].

To define a black hole, we consider a spacetime which has a connected future null infinity \mathcal{I}^+ . If \mathcal{I}^+ is not connected, it is sufficient to consider its connected components separately. The *causal past* of \mathcal{I}^+ , denoted $I^-(\mathcal{I}^+)$, is the region of \mathcal{M} from which there exist causal curves which reach \mathcal{I}^+ . A spacetime (\mathcal{M}, g) contains a *black hole* if $I^-(\mathcal{I}^+) \neq \mathcal{M}$. The black hole region is then defined as $\mathcal{B} = \mathcal{M} \setminus I^-(\mathcal{I}^+)$. The *event horizon* of the black hole is the boundary of the black hole: $\mathcal{H} = \partial\mathcal{B}$. (Analogous definitions for a white hole and a white hole horizon can be formulated by time inversion, i.e., replacing future null infinity with past null infinity, and the causal future with the causal past.)

2.4.3 Black hole spacetimes

Schwarzschild solution

The first exact solution to Einstein's vacuum field equations was uncovered by Schwarzschild in 1916 [24]. The solution describes the exterior spacetime of a spherically symmetric non-rotating gravitational source.

In spherical coordinates $\{t, r, \theta, \phi\}$, the Schwarzschild solution is described by the line element

$$ds^2 = g_{ab} dx^a dx^b = -f(r) dt^2 + \frac{dr^2}{f(r)} + r^2 (d\theta^2 + \sin^2 \theta d\phi^2), \quad f(r) = 1 - \frac{2M}{r}. \quad (2.84)$$

The parameter M , which is interpreted as the mass of the source, completely determines the solution. Far from the source ($r \rightarrow \infty$) the line element (2.84) approaches the

Minkowski line element expressed in standard spherical polar coordinates. Note that, in the case $M = 0$, the Schwarzschild spacetime reduces to that of Minkowski.

The domain $r > 2M$ is referred to as the *exterior region*. The *black hole region* of Schwarzschild spacetime is given by $0 < r \leq 2M$; the *event horizon* of the black hole is the null surface $r = 2M$. From the form of the metric (2.84), it would appear that $r = 2M$ is a singularity of the spacetime; however, this can be removed by transforming to an alternative coordinate system (e.g. Kruskal–Szekeres or Eddington–Finkelstein). Thus, $r = 2M$ is only a coordinate singularity. The Schwarzschild solution does, however, admit a physical singularity at $r = 0$. To see that this is the case, one may compute the Kretschmann invariant, which is given by $R_{abcd}R^{abcd} = \frac{48M^2}{r^6}$ in spherical coordinates. Clearly, this quantity diverges as $r \rightarrow 0$.

The metric (2.84) is independent of t , so ∂_t is a Killing vector. This is timelike for $r > 2M$, so the exterior Schwarzschild spacetime is (at least) stationary. In fact, ∂_t is hypersurface-orthogonal, so the exterior Schwarzschild spacetime is *static*. The metric (2.84) is also independent of ϕ , so ∂_ϕ is a Killing vector. This is spacelike in the exterior region. There are two further spacelike Killing vectors, given by $\sin \phi \partial_\theta + \cot \theta \cos \phi \partial_\phi$ and $\cos \phi \partial_\theta - \cot \theta \sin \phi \partial_\phi$. The three spacelike Killing vectors comprise the SO(3) rotations; they are precisely the Killing vectors admitted by the unit two-sphere S^2 . The Schwarzschild spacetime is therefore *spherically symmetric*.

Birkhoff [27] demonstrated that the Schwarzschild solution (2.84) is the *unique* spherically symmetric solution to Einstein's field equations in vacuum; this statement is known as *Birkhoff's theorem*.

Kerr solution

In 1963, almost half a century after the development of Einstein's field equations, Kerr discovered a solution to the vacuum field equations which describes a rotating black hole [35]. The exterior spacetime geometry may be described in Boyer–Lindquist coordinates $\{t, r, \theta, \phi\}$ [46] by the line element

$$ds^2 = g_{ab} dx^a dx^b \tag{2.85}$$

$$= - \left(1 - \frac{2Mr}{\Sigma} \right) dt^2 + \frac{\Sigma}{\Delta} dr^2 + \Sigma d\theta^2 + \left(r^2 + a^2 + \frac{2Ma^2r}{\Sigma} \sin^2 \theta \right) \sin^2 \theta d\phi^2 - \frac{4Mar \sin^2 \theta}{\Sigma} dt d\phi, \tag{2.86}$$

where $\Sigma(r, \theta) = r^2 + a^2 \cos^2 \theta$ and $\Delta(r) = r^2 - 2Mr + a^2$. Here, M and $J = aM$ are the mass and angular momentum of the black hole, respectively. The interpretation of these parameters is best understood by considering the asymptotic form of the metric (2.86), as shown in [87].

The Kerr spacetime has two event horizons at $r_{\pm} = M \pm \sqrt{M^2 - a^2 \cos^2 \theta}$. There exist two *ergosurfaces* at $r_{\pm}^{\text{ergo}} = M \pm \sqrt{M^2 - a^2 \cos^2 \theta}$. There is a physical curvature singularity (at which $R_{abcd}R^{abcd} \rightarrow \infty$) at $r = 0$ and $\theta = \frac{\pi}{2}$; this is in fact a ring singularity, which is best seen by transforming to Kerr–Schild coordinates [88]. For a comprehensive discussion of the key features of the Kerr spacetime, see e.g. [87, 89].

Clearly, the metric (2.86) is independent of the coordinates t and ϕ . In Boyer–Lindquist coordinates, the vectors $\xi_{(t)} = \partial_t$ and $\xi_{(\phi)} = \partial_\phi$ are therefore Killing vectors, which satisfy $\xi_{(a;b)} = 0$. These vectors correspond to time translation and rotation about the symmetry axis, respectively. In addition, any constant coefficient linear combination of these two Killing vectors is also a Killing vector. This exhausts the set of all Killing vectors of the Kerr spacetime. The spacetime is stationary (but not static), and axisymmetric. We note also that the metric (2.86) is invariant under simultaneous inversion of t and ϕ ($t \mapsto -t$ and $\phi \mapsto -\phi$).

Kerr–Newman solution

In 1965, Newman [36,37] discovered the solution to the Einstein–Maxwell field equations in *electrovacuum* which describes a rotating black hole endowed with an electric charge; this can be viewed as a generalisation of the Kerr metric. In Boyer–Lindquist coordinates, the line element is given by (2.86) with $\Sigma(r, \theta) = r^2 + a^2 \cos^2 \theta$ and $\Delta(r) = r^2 - 2Mr + a^2 + Q^2$. Here, M and $J = aM$ are again interpreted as the mass and angular momentum of the black hole, respectively; the parameter Q is interpreted as the black hole’s electric charge. Again, there exist a pair of event horizons, located at Boyer–Lindquist radii $r_\pm = M \pm \sqrt{M^2 - (a^2 + Q^2)}$. For $M^2 > a^2 + Q^2$, there are two distinct event horizons at r_\pm . In the case $M^2 = a^2 + Q^2$, the black hole is referred to as *extremal*, and the inner and outer horizons coincide at $r = M$. If $M^2 < a^2 + Q^2$, the Kerr–Newman solution describes a *naked singularity* [49].

As with the uncharged Kerr black hole, the Kerr–Newman solution is stationary and axisymmetric; these symmetries are encoded by the Killing vector fields $\xi_{(t)} = \partial_t$ and $\xi_{(\phi)} = \partial_\phi$. In Section 2.6.4, we will discuss the importance of these (and other) symmetries in the context of geodesic motion on the Kerr–Newman spacetime.

The Kerr–Newman family of solutions is characterised by three parameters: mass M ; spin a ; and electric charge Q . The uncharged ($Q = 0$) non-rotating ($a = 0$) case is simply the Schwarzschild solution, described above. The charged ($Q \neq 0$) non-rotating ($a = 0$) case is the Reissner–Nordström solution, which is a static spherically symmetric solution to the electrovacuum field equations, discovered by Reissner [25], Nordström [26] and others [90,91] between 1916 and 1921. The uncharged ($Q = 0$) rotating ($a \neq 0$) case is the Kerr solution, described above.

It should be emphasised that, in the case of rotating spacetimes, there is no analogue of Birkhoff’s theorem: it is *not* true in general that the vacuum geometry in the exterior region of a rotating object is given by (part of) the Kerr geometry. However, there are a number of powerful uniqueness theorems in situations with less symmetry. The so-called “no-hair theorem” states that isolated (mathematical) black holes in stationary axisymmetric electrovacuum are remarkably simple objects, and are uniquely characterised by three (externally observable) parameters: mass M , angular momentum J , and electric charge Q [92]. (The latter is thought to be negligible for astrophysical black holes.) In other words, black holes are described by the Kerr(–Newman) solution to the Einstein(–Maxwell) equations.

2.5 Newman–Penrose formalism

The Newman–Penrose formalism [93] is a special case of the tetrad formalism [89], in which tensorial quantities in general relativity are projected onto a complex null tetrad. This formalism is well-adapted to spacetime symmetries, and is a useful tool in the treatment of radiation on curved spacetime. In this section, we present a brief review of the tetrad formalism in general relativity, before outlining the key features of the (tensorial) Newman–Penrose formalism, which are relevant for the purposes of this thesis. A more exhaustive review of the Newman–Penrose formalism is presented by Stephani *et al.* [83], using both tensor and spinor notation.

2.5.1 Tetrad formalism

Tetrad representation

At each point in spacetime, one may construct a local basis of four contravariant vectors, i.e., a *tetrad* $\{e_{(\alpha)}{}^a\}$ [89]. Here, lowercase Greek letters enclosed in parentheses denote *tetrad indices*, whereas lowercase Latin letters are reserved for *tensor indices*. To each contravariant vector, we associate a covariant vector

$$e_{(\alpha)a} = g_{ab}e_{(\alpha)}{}^b, \quad (2.87)$$

so that spacetime indices are lowered (raised) with the spacetime metric (inverse metric). The inverse of $e_{(\alpha)}{}^a$, denoted $e^{(\alpha)}{}_a$, is defined such that $e_{(\alpha)}{}^a e^{(\beta)}{}_a = \delta_{(\alpha)}^{(\beta)}$ and $e_{(\alpha)}{}^a e^{(\alpha)}{}_b = \delta^a_b$, where the summation over tetrad/spacetime indices is assumed. Furthermore, we assume that

$$e_{(\alpha)}{}^a e_{(\beta)a} = S_{(\alpha)(\beta)}, \quad (2.88)$$

where $S_{(\alpha)(\beta)}$ is a symmetric matrix, often referred to as the *inner product matrix*. When all of the inner products between the tetrad legs (i.e., all components of $S_{(\alpha)(\beta)}$) are constant, the tetrad is referred to as a *rigid frame*. Let $S^{(\alpha)(\beta)}$ be the inverse of $S_{(\alpha)(\beta)}$, so that $S^{(\alpha)(\gamma)}S_{(\gamma)(\beta)} = \delta^{(\alpha)}_{(\beta)}$. Combining the above definitions, we see that

$$S_{(\alpha)(\beta)} e^{(\alpha)}{}_a = e_{(\beta)a}, \quad S^{(\alpha)(\beta)} e_{(\alpha)a} = e^{(\beta)}{}_a, \quad (2.89)$$

so that tetrad indices are lowered (raised) with the inner product matrix (inverse inner product matrix). Importantly, we also find

$$e_{(\alpha)a}e^{(\alpha)}{}_b = g_{ab}, \quad e_{(\alpha)}{}^a e^{(\alpha)b} = g^{ab}, \quad (2.90)$$

so the metric and its inverse can be expressed as an inner product of the tetrad legs.

Given any type- (r, s) tensor field with components $T^{a_1 \dots a_r}_{b_1 \dots b_s}$, we can project it onto the tetrad to obtain the *tetrad components*, viz.

$$T^{(\alpha_1) \dots (\alpha_r)}_{(\beta_1) \dots (\beta_s)} = e^{(\alpha_1)}{}_{a_1} \dots e^{(\alpha_r)}{}_{a_r} e_{(\beta_1)}{}^{b_1} \dots e_{(\beta_s)}{}^{b_s} T^{a_1 \dots a_r}_{b_1 \dots b_s}. \quad (2.91)$$

Similarly, we may pass from tetrad indices to spacetime indices according to

$$T^{a_1 \dots a_r}_{b_1 \dots b_s} = e_{(\alpha_1)}{}^{a_1} \dots e_{(\alpha_r)}{}^{a_r} e^{(\beta_1)}{}_{b_1} \dots e^{(\beta_s)}{}_{b_s} T^{(\alpha_1) \dots (\alpha_r)}_{(\beta_1) \dots (\beta_s)}. \quad (2.92)$$

Ricci rotation coefficients

The covariant derivative of a tetrad leg $e_{(\beta)}$ in the direction of another tetrad leg $e_{(\alpha)}$ can be expanded in the tetrad in terms of connection coefficients (see Section 2.1.4) as follows [83, 89]:

$$e_{(\alpha)}{}^b \nabla_b e_{(\beta)}{}^a = \Gamma_{(\alpha)(\beta)}^{(\gamma)} e_{(\gamma)}{}^a. \quad (2.93)$$

Contracting this equation with the basis vector $e^{(\delta)}{}_a$, using the properties of the tetrad listed above, and relabelling the indices, we arrive at the expression

$$\Gamma_{(\alpha)(\beta)(\gamma)} = e_{(\alpha)}{}^a e_{(\beta)a;b} e_{(\gamma)}{}^b. \quad (2.94)$$

The quantities $\Gamma_{(\alpha)(\beta)(\gamma)}$ are called the *Ricci rotation coefficients*.

If $\{e_{(\alpha)}\}$ is a rigid frame, such that the inner products $g_{ab} e_{(\alpha)}{}^a e_{(\beta)}{}^b = S_{(\alpha)(\beta)}$ are all constant, then the Ricci rotation coefficients are antisymmetric in the first pair of indices, i.e., $\Gamma_{(\alpha)(\beta)(\gamma)} = -\Gamma_{(\beta)(\alpha)(\gamma)}$.

Ricci and Bianchi identities

The Ricci identity (2.39) for the tetrad $\{e_{(\alpha)}\}$ is

$$e_{(\alpha)a;bc} - e_{(\alpha)a;cb} = R_{dabc} e_{(\alpha)}{}^d. \quad (2.95)$$

Projecting (2.95) onto the tetrad, and replacing the covariant derivatives of the tetrad legs using (2.94), we may express the projection of the Riemann tensor onto the tetrad in terms of the Ricci rotation coefficients as

$$\begin{aligned} R_{(\alpha)(\beta)(\gamma)(\delta)} &= \Gamma_{(\alpha)(\beta)(\delta);(\gamma)} - \Gamma_{(\alpha)(\beta)(\gamma);(\delta)} + \Gamma_{(\alpha)(\beta)(\eta)} \Gamma_{(\delta)}^{(\eta)}{}_{(\gamma)} \\ &\quad - \Gamma_{(\alpha)(\beta)(\eta)} \Gamma_{(\gamma)}^{(\eta)}{}_{(\delta)} + \Gamma_{(\eta)(\alpha)(\gamma)} \Gamma_{(\beta)}^{(\eta)}{}_{(\delta)} - \Gamma_{(\eta)(\alpha)(\delta)} \Gamma_{(\beta)}^{(\eta)}{}_{(\gamma)}, \end{aligned} \quad (2.96)$$

where we have used the notation $\Gamma_{(\alpha)(\beta)(\delta);(\gamma)} = e_{(\gamma)}{}^c \Gamma_{(\alpha)(\beta)(\delta);c}$.

Similarly, one may project the second (differential) Bianchi identity (2.42) onto the tetrad $\{e_{(\alpha)}\}$. The tetrad components of the covariant derivative of the Riemann tensor are

$$\begin{aligned} R_{(\alpha)(\beta)(\gamma)(\delta);(\eta)} &= R_{(\alpha)(\beta)(\gamma)(\delta),(\eta)} - \Gamma_{(\mu)(\alpha)(\eta)} R_{(\beta)(\gamma)(\delta)}^{(\mu)} - \Gamma_{(\mu)(\beta)(\eta)} R_{(\alpha)(\gamma)(\delta)}^{(\mu)} \\ &\quad - \Gamma_{(\mu)(\gamma)(\eta)} R_{(\alpha)(\beta)(\delta)}^{(\mu)} - \Gamma_{(\mu)(\delta)(\eta)} R_{(\alpha)(\beta)(\gamma)}^{(\mu)}, \end{aligned} \quad (2.97)$$

where $R_{(\alpha)(\beta)(\gamma)(\delta);(\eta)} = e_{(\eta)}{}^a R_{(\alpha)(\beta)(\gamma)(\delta);a}$ and $R_{(\alpha)(\beta)(\gamma)(\delta),(\eta)} = e_{(\eta)}{}^a R_{(\alpha)(\beta)(\gamma)(\delta),a}$. The projection of the second Bianchi identity (2.42) onto the tetrad is then given by anti-symmetrising (2.97) over its last three indices.

Orthonormal tetrads

A special case of the tetrad formalism involves choosing an orthonormal tetrad at each spacetime point. An *orthonormal tetrad* consists of one unit timelike vector and three unit spacelike vectors whose inner products satisfy

$$g_{ab} e_{(\alpha)}{}^a e_{(\beta)}{}^b = \eta_{(\alpha)(\beta)}, \quad (2.98)$$

where $\eta_{(\alpha)(\beta)}$ are the components of the Minkowski metric; this is a special case of (2.88). The spacetime metric and its inverse may be expressed in terms of the orthonormal tetrad, respectively, as

$$g_{ab} = -e^{(0)}_a e^{(0)}_b + e^{(1)}_a e^{(1)}_b + e^{(2)}_a e^{(2)}_b + e^{(3)}_a e^{(3)}_b, \quad (2.99)$$

$$g^{ab} = -e_{(0)}^a e_{(0)}^b + e_{(1)}^a e_{(1)}^b + e_{(2)}^a e_{(2)}^b + e_{(3)}^a e_{(3)}^b. \quad (2.100)$$

Complex null tetrads

In general relativity, complex null tetrads also play an important role. A *complex null tetrad* $\{k, n, m, \bar{m}\}$ consists of two real null vectors (k and n) and a pair of complex conjugate null vectors (m and \bar{m}), which satisfy

$$k^a n_a = -1, \quad m^a \bar{m}_a = 1, \quad (2.101)$$

with all other inner products zero. The metric and inverse metric can be written, respectively, in the form

$$g_{ab} = -k_a n_b - n_a k_b + m_a \bar{m}_b + \bar{m}_a m_b, \quad (2.102)$$

$$g^{ab} = -k^a n^b - n^a k^b + m^a \bar{m}^b + \bar{m}^a m^b. \quad (2.103)$$

One may always construct a complex null tetrad $\{k, n, m, \bar{m}\}$ from a real orthonormal tetrad $\{e_{(0)}, e_{(1)}, e_{(2)}, e_{(3)}\}$ as follows:

$$k = \frac{1}{\sqrt{2}} (e_{(0)} + e_{(3)}), \quad n = \frac{1}{\sqrt{2}} (e_{(0)} - e_{(3)}), \quad (2.104)$$

$$m = \frac{1}{\sqrt{2}} (e_{(1)} + ie_{(2)}), \quad \bar{m} = \frac{1}{\sqrt{2}} (e_{(1)} - ie_{(2)}). \quad (2.105)$$

It is standard to introduce four directional derivatives along each of the legs of the complex null tetrad,

$$D = k^a \nabla_a, \quad \Delta = n^a \nabla_a, \quad \delta = m^a \nabla_a, \quad \bar{\delta} = \bar{m}^a \nabla_a. \quad (2.106)$$

2.5.2 Transformation laws for complex null tetrads

A complex null tetrad $\{k, n, m, \bar{m}\}$ may be transformed in the following ways [83].

(i) Null rotations which keep k fixed:

$$k' = k, \quad m' = m + Bk, \quad n' = n + B\bar{m} + \bar{B}m + B\bar{B}k, \quad (2.107)$$

where B is a complex function.

(ii) Null rotations which keep n fixed:

$$n' = n, \quad m' = m + En, \quad k' = k + E\bar{m} + \bar{E}m + E\bar{E}n, \quad (2.108)$$

where E is a complex function.

(iii) Spatial rotations in the (m, \bar{m}) -plane:

$$m' = e^{i\Theta} m, \quad (2.109)$$

where Θ is a real function.

(iv) Boosts in the (k, n) -plane:

$$k' = Ak, \quad n' = \frac{1}{A}n, \quad (2.110)$$

where $A > 0$ is a real function.

Together, the transformations (2.107)–(2.110), which contain six real parameters, represent the six-parameter group of Lorentz transformations.

2.5.3 Spin coefficients

Newman and Penrose [93] introduce a set of twelve independent complex scalars, which are defined as linear combinations of the Ricci rotation coefficients (2.94) for the complex null tetrad $\{k, n, m, \bar{m}\}$. The *spin coefficients* (or *Newman–Penrose scalars*) are [83, 89]

$$\kappa = -\Gamma_{abc} m^a k^b k^c = -m^a k_{a;b} k^b, \quad \rho = -\Gamma_{abc} m^a k^b \bar{m}^c = -m^a k_{a;b} \bar{m}^b, \quad (2.111)$$

$$\sigma = -\Gamma_{abc} m^a k^b m^c = -m^a k_{a;b} m^b, \quad \tau = -\Gamma_{abc} m^a k^b n^c = -m^a k_{a;b} n^b, \quad (2.112)$$

$$\pi = \Gamma_{abc} \bar{m}^a n^b k^c = \bar{m}^a n_{a;b} k^b, \quad \nu = \Gamma_{abc} \bar{m}^a n^b n^c = \bar{m}^a n_{a;b} n^b, \quad (2.113)$$

$$\mu = \Gamma_{abc} \bar{m}^a n^b m^c = \bar{m}^a n_{a;b} m^b, \quad \lambda = \Gamma_{abc} \bar{m}^a n^b \bar{m}^c = \bar{m}^a n_{a;b} \bar{m}^b, \quad (2.114)$$

$$\begin{aligned} \varepsilon &= \frac{1}{2} \left(\Gamma_{abc} \bar{m}^a m^b k^c - \Gamma_{abc} n^a k^b k^c \right) & \gamma &= \frac{1}{2} \left(\Gamma_{abc} k^a n^b k^c - \Gamma_{abc} m^a \bar{m}^b n^c \right) \\ &= \frac{1}{2} \left(\bar{m}^a m_{a;b} k^b - n^a k_{a;b} k^b \right), & &= \frac{1}{2} \left(k^a n_{a;b} k^b - m^a \bar{m}_{a;b} n^b \right), \end{aligned} \quad (2.115)$$

$$\begin{aligned} \alpha &= \frac{1}{2} \left(\Gamma_{abc} k^a n^b \bar{m}^c - \Gamma_{abc} \bar{m}^a m^b \bar{m}^c \right) & \beta &= \frac{1}{2} \left(\Gamma_{abc} \bar{m}^a m^b m^c - \Gamma_{abc} n^a k^b m^c \right) \\ &= \frac{1}{2} \left(k^a n_{a;b} \bar{m}^b - \bar{m}^a \bar{m}_{a;b} m^b \right), & &= \frac{1}{2} \left(\bar{m}^a m_{a;b} m^b - n^a k_{a;b} m^b \right). \end{aligned} \quad (2.116)$$

The twelve spin coefficients are the primary quantities in the Newman–Penrose formalism. For this reason, the formalism is sometimes referred to in the literature as the *spin coefficient formalism*.

In many situations, the complex null tetrad will be transformed using one of the Lorentz transformations (i)–(iv) from Section 2.5.2. The corresponding transformation laws for the Newman–Penrose scalars can be derived using the definitions involving projections of the tetrad legs (2.111)–(2.116); alternatively, these transformation laws can be found in Chapter 7 of [83].

The directional derivatives (2.106) satisfy a set of commutation relations involving the Newman–Penrose scalars, which arise from the metric compatibility and torsion-free property of the covariant derivative; the commutator identities are

$$\Delta D - D\Delta = (\gamma + \bar{\gamma})D + (\varepsilon + \bar{\varepsilon})\Delta - (\bar{\tau} + \pi)\delta - (\tau + \bar{\pi})\bar{\delta}, \quad (2.117)$$

$$\delta D - D\delta = (\bar{\alpha} + \beta - \bar{\pi})D + \kappa\Delta - (\bar{\rho} + \varepsilon - \bar{\varepsilon})\delta - \sigma\bar{\delta}, \quad (2.118)$$

$$\delta\Delta - \Delta\delta = -\bar{\nu}D + (\tau - \bar{\alpha} - \beta)\Delta + (\mu - \gamma + \bar{\gamma})\delta + \bar{\lambda}\bar{\delta}, \quad (2.119)$$

$$\bar{\delta}\delta - \delta\bar{\delta} = (\bar{\mu} - \mu)D + (\bar{\rho} - \rho)\Delta + (\alpha - \bar{\beta})\delta - (\bar{\alpha} - \beta)\bar{\delta}. \quad (2.120)$$

2.5.4 Weyl scalars

In Section 2.1.8, we introduced the Weyl tensor C_{abcd} , the trace-free part of the Riemann curvature tensor; see (2.45). In four spacetime dimensions, the Weyl tensor has ten independent components, which are encoded in five complex scalars in the Newman–Penrose formalism. The *Weyl scalars* are

$$\Psi_0 = C_{abcd}k^ak^bk^cm^d, \quad (2.121)$$

$$\Psi_1 = C_{abcd}k^an^bk^cm^d, \quad (2.122)$$

$$\Psi_2 = C_{abcd}k^am^b\bar{m}^cn^d, \quad (2.123)$$

$$\Psi_3 = C_{abcd}k^an^b\bar{m}^cn^d, \quad (2.124)$$

$$\Psi_4 = C_{abcd}\bar{m}^an^b\bar{m}^cn^d. \quad (2.125)$$

The Ricci tensor R_{ab} has ten independent components. These are encoded in the Ricci curvature scalar R , and the six complex *Ricci scalars* and their complex conjugates:

$$\Phi_{00} = \frac{1}{2}R_{ab}k^ak^b = \bar{\Phi}_{00}, \quad \Phi_{01} = \frac{1}{2}R_{ab}k^am^b = \bar{\Phi}_{10}, \quad (2.126)$$

$$\Phi_{02} = \frac{1}{2}R_{ab}m^am^b = \bar{\Phi}_{20}, \quad \Phi_{11} = \frac{1}{2}R_{ab}(k^an^b + m^a\bar{m}^b) = \bar{\Phi}_{11}, \quad (2.127)$$

$$\Phi_{12} = \frac{1}{2}R_{ab}n^am^b = \bar{\Phi}_{21}, \quad \Phi_{22} = \frac{1}{2}R_{ab}n^an^b = \bar{\Phi}_{22}. \quad (2.128)$$

We note that there are nine independent components of the complex Ricci scalars (2.126)–(2.128). In these definitions, one may replace the Ricci tensor R_{ab} with its trace-free counterpart $S_{ab} = R_{ab} - \frac{1}{4}Rg_{ab}$ if so desired.

2.5.5 Newman–Penrose field equations

When projected onto the complex null tetrad $\{k, n, m, \bar{m}\}$, the Ricci identities (2.96) give rise to a system of 36 equations involving the Newman–Penrose scalars (2.111)–(2.116), their directional derivatives along tetrad legs (2.106), and the Weyl and Ricci scalars (2.121)–(2.128). The full set of *Ricci equations* (or *Newman–Penrose field equations*) are written out explicitly in Chapter 7 of [83].

In addition, one may project the Bianchi identities onto the complex null tetrad $\{k, n, m, \bar{m}\}$, which can be achieved by anti-symmetrising the expression (2.97) over its last three indices. This yields a system of equations involving the Newman–Penrose scalars, the Weyl and Ricci scalars and their directional derivatives. The full set of such equations, referred to as the *Bianchi equations*, can again be found in Chapter 7 of [83].

2.5.6 Geroch–Held–Penrose formalism

Geroch *et al.* [94] developed a formalism for the treatment of general relativity that lies between a fully covariant formulation and the Newman–Penrose spin-coefficient formalism; this approach is referred to as the *Geroch–Held–Penrose formalism*. In this approach, only a *pair* of null directions – as opposed to an entire null tetrad – is considered at each spacetime point. The formulae in this Geroch–Held–Penrose formalism are simpler than those of the spin-coefficient formalism of Newman and Penrose.

The most general transformation which preserves the pair of null directions is a two-parameter subgroup of the full group of Lorentz transformations, corresponding to a rotation in the (m, \bar{m}) -plane and a boost in the (k, n) -plane; see transformations (iii) and (iv) of Section 2.5.2. Under this transformation, the tetrad legs transform according to

$$k' = Ak, \quad n' = \frac{1}{A}n, \quad m' = e^{i\Theta}m, \quad \bar{m}' = e^{-i\Theta}\bar{m}, \quad (2.129)$$

where $A = C\bar{C}$ and $e^{i\Theta} = C\bar{C}^{-1}$, with C complex.

A scalar quantity z which transforms according to $z' = C^p\bar{C}^qz$ under (2.129) is called a *scalar of type* (p, q) . The scalar is said to have *spin weight* $\mathfrak{s} = \frac{1}{2}(p - q)$, and *boost weight* $\mathfrak{b} = \frac{1}{2}(p + q)$. The spin (boost) weight can be obtained easily by counting the number of times the legs m and \bar{m} (k and n) appear in the definition of the quantity: we add 1 for the number of times m (k) occurs, and subtract 1 for the number of times \bar{m} (n) occurs. For example, the Newman–Penrose scalar $\rho = -m^ak_{a;b}\bar{m}^b$ has spin weight $\mathfrak{s} = 0$, since there is one m and one \bar{m} in the definition of ρ ; and it has boost weight $\mathfrak{b} = 1$, since k appears once and n does not appear. We note here that certain quantities (e.g. α) do not have a well-defined spin/boost weight.

2.6 Integrability and chaos in general relativity

The main theme of this thesis is the study of light propagation on black hole spacetimes. As described in Section 2.2, light follows null geodesics on curved spacetime; this can be modelled using a Hamiltonian formalism, where the geodesics are integral curves of Hamilton's equations with Hamiltonian function $H = \frac{1}{2}g^{ab}p_ap_b$. The themes of integrability and chaos in Hamiltonian dynamical systems will be of central importance to our understanding of light propagation on black hole spacetimes in general relativity.

In this section, we review some key features of the theory of (non-linear) dynamical systems. In particular, we will discuss the relationship between spacetime symmetries, conserved quantities along geodesics, and integrability in Section 2.6.4. We will also present a review of some important features of the theory of deterministic chaos in Section 2.6.5.

2.6.1 Dynamical systems

Dynamical systems theory is concerned with the long-term behaviour of evolving systems. In general, a *dynamical system* consists of a *state space* (or *phase space*), a space whose points describe the state of the system at any instant of time; and a rule which governs the time-evolution of points in the state space. For a comprehensive overview of this topic, see e.g. [95–100]. Here, we review some key concepts and definitions.

Mathematically, a *discrete-time dynamical system* consists of a non-empty set X (the state space) and a map $f: X \rightarrow X$ (the time-evolution rule). For $n \in \mathbb{N}$, the n th iterate of f is given by $f^n = f \circ \dots \circ f$, the n -fold composition of f , where f^0 is the identity map. If f is invertible, then $f^{-n} = f^{-1} \circ \dots \circ f^{-1}$, i.e., the inverse of f composed with itself n times. The iterates satisfy $f^{n+m} = f^n \circ f^m$, so they form a group if f is invertible. In this context, a discrete-time dynamical system is often referred to as a *map*.

Typically, the set X will be endowed with some additional structure which is preserved by the map f . For example, (X, f) could be a smooth manifold and a smooth map (or a diffeomorphism in the case of invertible maps).

A *continuous time dynamical system* consists of a space X and a one-parameter family of maps $\{f^\lambda: X \rightarrow X \mid \lambda \in \mathbb{R}\}$ which form a one-parameter group, i.e., $f^{\lambda+\tau} = f^\lambda \circ f^\tau$ and f^0 is the identity map. Here, λ is the continuous time parameter. In this context, the dynamical system is called a *flow*. The map is invertible, since $(f^\lambda)^{-1} = f^{-\lambda}$. For a fixed value $\lambda = \lambda_0$, the iterates $(f^{\lambda_0})^n$ form a discrete-time dynamical system.

The time-evolution of many systems of interest is governed by systems of ordinary differential equations. If there is no time-dependent forcing, then the system will be *autonomous*. We consider a continuous-time dynamical system described by a set of first-order ordinary differential equations

$$\dot{z}(\lambda) = X(z), \quad (2.130)$$

where X is a vector field, $z \in \mathcal{P}$ represents a point in phase space \mathcal{P} , and an overdot denotes differentiation with respect to the time parameter λ . The system (2.130) induces a flow

$$z = \varphi_X(z_0, \lambda), \quad (2.131)$$

which is a non-linear function of λ and the initial data z_0 . Systems of higher-order ordinary differential equations can, in general, be reduced to the form (2.130). Phase space volumes may be conserved if the vector field F is divergence-free; otherwise, phase space volumes may expand or contract. Dynamical systems are divided into two main groups: Hamiltonian (conservative) systems and dissipative systems. In this work we focus on the former.

2.6.2 Hamiltonian systems

A particularly special class of dynamical system is a Hamiltonian system. In the Hamiltonian approach, a dynamical system is described in terms of phase space, which can be represented geometrically using symplectic geometry. For a more detailed description of Hamiltonian systems, see the textbooks by Arnold [84], Ozorio de Almeida [101], and Lowenstein [102], for example.

Let \mathcal{P} be a $2N$ -dimensional manifold. A *symplectic structure* on \mathcal{P} is a non-degenerate two-form Ω which is closed ($d\Omega = 0$). The pair (\mathcal{P}, Ω) is called a *symplectic manifold*; it describes a dynamical system with N degrees of freedom. The existence of a non-degenerate two-form Ω implies that the manifold \mathcal{P} must be even-dimensional.

Let z^A ($A \in \{1, \dots, 2N\}$) denote coordinates on \mathcal{P} . The components of the symplectic form Ω_{AB} are given by an anti-symmetric non-degenerate matrix. The inverse symplectic form, whose components are written Ω^{AB} , is defined via the relation $\Omega_{AC}\Omega^{BC} = \delta_A^B$.

An *observable* is a scalar function on \mathcal{P} . In autonomous Hamiltonian systems, the observables do not explicitly depend on time. For an observable F , the symplectic form defines the corresponding *Hamiltonian vector field* X_F , which is given in components by

$$X_F^A = \Omega^{AB} F_{,B}, \quad (2.132)$$

where a comma denotes the partial derivative with respect to the phase space coordinate z^A , i.e., $F_{,A} = \frac{\partial F}{\partial z^A}$. The integral curves of the vector field X_F determine a map from \mathcal{P} to itself; this is called a *Hamiltonian flow*. Introducing a parameter λ and local coordinates $z^A(\lambda)$, the integral curves of the Hamiltonian flow are given by

$$X_F = \frac{dz^A}{d\lambda} \partial_A = \dot{z}^A \partial_A. \quad (2.133)$$

For a pair of observables F and G , the symplectic structure Ω gives another observable $K = \{F, G\}$, which is the *Poisson bracket*, given by

$$\{F, G\} = \Omega^{AB} F_{,A} G_{,B}. \quad (2.134)$$

The Poisson bracket satisfies the following properties.

- (i) *Antisymmetry*: $\{F, G\} = -\{G, F\}$ for all observables F and G on \mathcal{P} . This follows from the fact that Ω is a two-form
- (ii) *Jacobi identity*: $\{F, \{G, K\}\} + \{G, \{K, F\}\} + \{K, \{F, G\}\} = 0$, for all observables F , G and K on \mathcal{P} . This follows from the fact that Ω is closed.

A pair of observables F and G are said to *Poisson commute* if their Poisson bracket vanishes, i.e., $\{F, G\} = 0$. Such observables are said to be *in involution*. These properties ensure that observables on phase space form a Lie algebra with respect to the Poisson bracket. This is related to the Lie algebra of (Hamiltonian) vector fields via

$$[X_F, X_G] = -X_{\{F, G\}}, \quad (2.135)$$

where $[X, Y]$ denotes the commutator of the vector fields X and Y (see Section 2.1.5). If a pair of observables are in involution, then their corresponding Hamiltonian vector fields commute.

Hamiltonian vector fields preserve the symplectic two-form Ω , that is, $\mathcal{L}_{X_F} \Omega = 0$. Using the Leibniz rule for the Lie derivative of the exterior product, one may demonstrate that the volume form $\Omega^{\wedge N} = \Omega \wedge \dots \wedge \Omega$ (N times) induced by the symplectic structure is preserved by the Hamiltonian vector field X_F ; in other words, $\mathcal{L}_{X_F} \Omega^{\wedge N} = 0$. This is the famous *Liouville's theorem*, which says that the natural volume form on a symplectic manifold is invariant under the Hamiltonian flow.

The symplectic structure ensures the existence of a special class of coordinates, which simplify the expressions discussed above. The existence of such coordinates follows from the *Darboux theorem*, which states that, in the vicinity of a point of \mathcal{P} , it is possible to choose *canonical coordinates* $z^A = (q^a, p_b) = (q^1, \dots, q^N, p_1, \dots, p_N)$ (with $a, b \in \{1, \dots, N\}$), in which the symplectic form Ω can be written in canonical form as

$$\Omega = dq^a \wedge dp_a. \quad (2.136)$$

In canonical coordinates, the Hamiltonian vector field takes the form

$$X_F = \frac{\partial F}{\partial p_a} \frac{\partial}{\partial q^a} - \frac{\partial F}{\partial q^a} \frac{\partial}{\partial p_a}, \quad (2.137)$$

and the Poisson bracket is

$$\{F, G\} = \frac{\partial F}{\partial q^a} \frac{\partial G}{\partial p_a} - \frac{\partial F}{\partial p_a} \frac{\partial G}{\partial q^a}. \quad (2.138)$$

The dynamics on phase space is determined by the *Hamiltonian function* H , a given scalar function on \mathcal{P} . (For autonomous systems, H does not explicitly depend on time.) The time-evolution is determined by the Hamiltonian flow corresponding to the Hamiltonian function: phase space trajectories are the integral curves of the Hamiltonian vector field corresponding to H , given in canonical coordinates by

$$X_H = \dot{z}^A \partial_A = \dot{q}^a \frac{\partial}{\partial q^a} + \dot{p}_a \frac{\partial}{\partial p_a}. \quad (2.139)$$

Comparison of (2.137) with (2.139) yields *Hamilton's equations*,

$$\dot{q}^a = \frac{\partial H}{\partial p_a}, \quad \dot{p}_a = -\frac{\partial H}{\partial q^a}. \quad (2.140)$$

The time-evolution of an observable F on \mathcal{P} is given by

$$\dot{F} = \{F, H\}. \quad (2.141)$$

In the case of geodesic motion on a spacetime (\mathcal{M}, g) , the phase space is the cotangent bundle $\mathcal{P} = T^*\mathcal{M}$, and the canonical coordinates on \mathcal{P} are the spacetime coordinates $\{x^a\}$ and their conjugate momenta $\{p_a\}$ (see Section 2.2.3). In this work, we are concerned only with the study of null geodesic motion on four-dimensional spacetime (\mathcal{M}, g_{ab}) , so the phase space \mathcal{P} is eight-dimensional and there are four degrees of freedom. We will use either q^a or x^a to denote spacetime coordinates, depending on the context.

2.6.3 Integrability in Hamiltonian systems

A *conserved quantity* – also referred to as a *constant* or *integral of motion* – is an observable which remains constant along phase space trajectories. Clearly, an observable F is a conserved quantity if it Poisson commutes with the Hamiltonian, i.e., $\{F, H\} = 0$.

Each observable F on phase space induces a Hamiltonian flow, which can be viewed as a transformation on phase space. If the observable is a constant of motion, then the corresponding Hamiltonian vector field commutes with the time-evolution X_H : $[X_F, X_H] = 0$, by (2.135). Any trajectory which satisfies Hamilton's equations (2.140) can be mapped to another trajectory by the transformation induced by F . This says that conserved quantities generate symmetries of the time-evolution of a dynamical system. Frolov *et al.* [50] note this in the following theorem, which can be viewed as *Noether's theorem* on phase space. Let Y be a vector field which preserves both the symplectic structure and the Hamiltonian, i.e., $\mathcal{L}_Y \Omega = 0$ and $\mathcal{L}_Y H = 0$. Then there exists a constant of motion I , such that $Y = X_I$.

Of particular interest are dynamical systems which admit multiple observables which mutually Poisson commute (i.e., they are in involution). A dynamical system which admits the maximum possible number of independent commuting observables is said to be *completely integrable*. The time-evolution of such systems is “regular”: trajectories remain

in well-defined submanifolds of phase space, and can be obtained using a systematic procedure. When the system is non completely integrable, it is called *non-integrable*. In such systems, we generically observe rich chaotic motion in phase space; see Section 2.6.5.

A Hamiltonian system with N degrees of freedom is said to be *Liouville integrable* if it admits N independent constants of motion F_a which are in involution with one another, $\{F_a, F_b\} = 0$ (for all $a, b \in \{1, 2, \dots, N\}$); one of the F_a is the Hamiltonian function itself, say $H = F_1$. The fact that the F_a are independent means that no one of these integrals of motion can be expressed as a function of the other $N - 1$ conserved quantities. The requirement that an integrable system has N independent integrals of motion implies that the dynamical trajectories lie on N -dimensional surface defined by the equations

$$F_a(p_1, \dots, p_N, q_1, \dots, q_N) = C_a, \quad (2.142)$$

where C_a are constants given by the values of the phase space functions F_a along the dynamical trajectories. The requirement that the constants of motion be in involution restricts the topology of the N -dimensional surface (2.142): it must be an N -dimensional torus [84, 101].

The N -tori in a completely integrable system with N degrees of freedom in the $2N$ -dimensional phase space are often referred to as *invariant tori* [103]. This is due to the fact that any dynamical trajectory which starts on one of the N -tori remains on it forever. Different initial conditions lead to nested N -tori in phase space which do not intersect. The motion on these N -tori is characterised by a set of N fundamental frequencies ω_i , $i \in \{1, 2, \dots, N\}$. The trajectory then has n associated fundamental periods, given by $\frac{2\pi}{\omega_i}$. If the trajectory is *closed* on the N -torus, then it is *exactly periodic*. In order for this to be the case, the ratio of every pair of frequencies must be a rational number, i.e., $\frac{\omega_i}{\omega_j} \in \mathbb{Q}$ for all $i, j \in \{1, 2, \dots, N\}$. In such circumstances, the N -torus is called *resonant*. Conversely, if the frequency ratio $\frac{\omega_i}{\omega_j}$ is irrational, then the orbit will not be closed on the N -torus. In this case, we have *quasi-periodicity*, and the orbit is said to be *ergodic* on the N -torus [103].

Many completely integrable systems may be solved analytically by quadrature (i.e., using a finite number of algebraic operations and integrations). This prescription, which involves transforming to *action-angle variables*, is codified in *Liouville's procedure*, which is summarised in Appendix B of [50].

2.6.4 Killing objects, symmetries and integrability

Symmetries play a key role in general relativity, both in the search for cherished exact solutions to Einstein's field equations, and in the study of the integrability properties of wave equations and geodesic motion on a fixed spacetime background. Here, we briefly review the relationship between Killing vectors and Killing tensors (referred to collectively as Killing objects), symmetries on spacetime and phase space, and complete integrability of geodesic motion. A highly comprehensive account of this topic can be found in the review by Frolov *et al.* [50], and the review by Cariglia [104].

As discussed in Section 2.2, the motion of a free particle on a spacetime (\mathcal{M}, g_{ab}) follows a causal geodesic $x^a(\lambda)$, which is completely determined by the spacetime geometry, i.e.,

by the Hamiltonian $H = \frac{1}{2}g^{ab}p_ap_b$, where $p_a = g_{ab}\dot{x}^a$ are the canonical momenta, and an overdot denotes differentiation with respect to an affine parameter λ .

Killing vectors and explicit symmetries

We saw in Section 2.1.6 that Killing vector fields generate isometries of spacetime. According to Noether's theorem, there is a one-to-one correspondence between these isometries and constants of motion along geodesics. The latter can be expressed in terms of the Killing vector ξ^a and the momentum p_a as $I = \xi^a p_a$. It is quick to demonstrate that I is indeed an integral of motion:

$$\dot{I} = p^b I_{;b} = p^b (\xi^a p_a)_{;b} = p^b \xi_{a;b} p^a + p^b p_{a;b} \xi^a = \xi_{(a;b)} p^a p^b = 0, \quad (2.143)$$

where we have used the Killing vector equation $\xi_{(a;b)} = 0$, and the geodesic equation $p^b p_{a;b} = 0$ in the final step. On phase space \mathcal{P} , the Hamiltonian vector field corresponding to the constant of motion I is given by [50]

$$X_I = \xi^a \frac{\partial}{\partial x^a} - \xi^b_{;a} p_b \frac{\partial}{\partial p_a}. \quad (2.144)$$

When projected back onto the spacetime manifold \mathcal{M} , the Hamiltonian vector field reduces to the Killing vector field ξ . If the canonical projection of a phase space symmetry yields a well-defined spacetime symmetry, the symmetry is said to be *explicit*. Killing vectors therefore correspond to explicit symmetries. The projection of the phase space symmetry must depend only on the spacetime coordinates. It follows that the conserved quantity I must be linear in the momentum p_a . Thus, the constants of motion along a particle's trajectories in spacetime which correspond to explicit symmetries must be linear in the particle's momentum [50].

Killing tensors and hidden symmetries

In addition to explicit symmetries on spacetime induced by Killing vectors, there may exist symmetries which do not have a simple description on spacetime. Such symmetries are referred to as *hidden* symmetries, and are generated by higher-rank Killing tensors. A *Killing tensor* $K^{a_1 \dots a_r}$ of type $(r, 0)$ is a totally symmetric tensor which satisfies the *Killing tensor equation*,

$$\nabla^{(b} K^{a_1 \dots a_r)} = 0. \quad (2.145)$$

A type-(1, 0) Killing tensor is a Killing vector. A trivial example of a type-(2, 0) Killing tensor which exists on every spacetime is the (inverse) metric tensor g^{ab} ; this fact follows from the metric compatibility of the covariant derivative, i.e., $g_{ab;c} = 0$.

Using the Killing tensor equation (2.145) and the geodesic equation, one may show that, if $K^{a_1 \dots a_r}$ is a type- $(r, 0)$ Killing tensor, the scalar quantity

$$K = K^{a_1 \dots a_r} p_{a_1} \dots p_{a_r} \quad (2.146)$$

is conserved along geodesics. A special case of this result is the so-called Hamiltonian constraint (see Section 2.2): the Hamiltonian function $H = \frac{1}{2}g^{ab}p_ap_b$ is constant along dynamical trajectories (i.e., the geodesics), due to the fact that g^{ab} is a (trivial) Killing tensor.

Conformal Killing vectors, conformal Killing tensors and null geodesics

Our primary focus in this thesis is the propagation of electromagnetic radiation – in particular, light rays – on curved spacetime. We shall therefore now turn our attention to generalisations of Killing vectors and Killing tensors which provide conserved quantities only along *null* geodesics. A *conformal Killing vector* is a vector ξ which satisfies

$$\nabla^{(a}\xi^{b)} = \alpha g^{ab}, \quad (2.147)$$

for some function α . This is known as the *conformal Killing vector equation*. Clearly, if $\alpha = 0$, (2.147) reduces to the Killing vector equation (2.33). Now, given a conformal Killing vector ξ , it is quick to show that the observable $I = \xi^a p_a$ is conserved along *null* geodesics (which have $g^{ab}p_a p_b = 0$):

$$\dot{I} = p^b I_{;b} = p^b (\xi^a p_a)_{;b} = p^b \xi_{a;b} p^a + p^b p_{a;b} \xi^a = \xi_{(a;b)} p^a p^b = \alpha g^{ab} p_a p_b = 0. \quad (2.148)$$

It is also possible to show that both Killing vectors and conformal Killing vectors yield conserved quantities for any matter fields whose stress–energy tensors T^{ab} are (i) divergence-free ($\nabla_a T^{ab} = 0$), and (ii) traceless ($T = T^a_a = 0$). In particular, the current $J^a = T^{ab} \xi_b$ is conserved: $\nabla_a J^a = 0$ [50].

A symmetric tensor $K^{a_1 \dots a_r}$ of type $(r, 0)$ is a *conformal Killing tensor* if it satisfies the *conformal Killing tensor equation*

$$\nabla^{(b} K^{a_1 \dots a_r)} = g^{(ba_1} \alpha^{a_2 \dots a_r)}, \quad (2.149)$$

where α is a symmetric tensor of type $(r - 1, 0)$. Of course, in the case $\alpha = 0$, (2.149) reduces to the Killing tensor equation (2.145). The observable $K = K^{a_1 \dots a_r} p_{a_1} \dots p_{a_r}$ is conserved along *null* geodesics if $K^{a_1 \dots a_r}$ is a conformal Killing tensor.

The principal tensor

We now turn our attention to a highly important geometrical object – the *principal tensor* – which is responsible for a complete set of explicit and hidden symmetries, and whose existence uniquely determines the spacetime geometry, which is given by the Kerr–NUT–(anti-)de Sitter metric [50]. The *principal tensor* h_{ab} is a non-degenerate closed conformal Killing–Yano two-form, which obeys

$$\nabla_c h_{ab} = g_{ca} \xi_b - g_{cb} \xi_a, \quad (2.150)$$

where ξ^a is an associated vector field. Since h is closed, there exists (locally) a potential one-form, say b , such that $h = db$. By contracting (2.150) with g^{cb} and using the antisymmetry of h_{ab} , one may show that, in four spacetime dimensions, the vector ξ^a takes the form

$$\xi_a = -\frac{1}{3} \nabla^b h_{ab}. \quad (2.151)$$

It is possible to show that [105–108]

$$\mathcal{L}_\xi g = 0; \quad (2.152)$$

in other words, ξ^a defined in (2.151) is a Killing vector; this is referred to as the *primary Killing vector*.

The existence of a principal tensor is responsible for some remarkable properties. The algebraic properties of the principal tensor may be exploited to construct a (*special*) *Darboux frame* [50]. Moreover, the principal tensor is responsible for the existence of a rich symmetry structure, referred to as the *Killing tower*. In particular, the principal tensor can be used to generate (i) a family of closed conformal Killing–Yano forms, generated from exterior powers of the principal tensor; (ii) a family of Killing–Yano forms, which are the Hodge duals of the closed conformal Killing–Yano forms; (iii) rank-two Killing tensors, which are the “squares” of the Killing–Yano forms; (iv) rank-two conformal Killing tensors, which are the squares of the closed conformal Killing–Yano forms; and (v) Killing vectors, which are given by contracting the Killing tensors with the primary Killing vector. For a more detailed description of this procedure, see Chapter 2 of the review by Frolov *et al.* [50].

The most general higher-dimensional black hole solution to Einstein’s field equations – the (on-shell) Kerr–NUT–(anti-)de Sitter metric – admits a principal tensor. This result was first proved for Myers–Perry black holes [109], and later generalised to the more general family of Kerr–NUT–(anti-)de Sitter spacetimes [110]. In fact, the most general vacuum solution with a cosmological constant to the Einstein field equations in any number of spacetime dimensions ($n \geq 4$) that admits a principal tensor is precisely the Kerr–NUT–(anti-)de Sitter metric [106, 107, 111].

Clearly, the very existence of a principal restricts the allowed geometry. In fact, the existence of this special geometrical object determines the geometry entirely: the most general geometry that is consistent with the existence of the principal tensor is the so-called *off-shell* Kerr–NUT–(anti-)de Sitter geometry; see [50] and references therein.

The role of the principal tensor on Kerr spacetime

Let us now illustrate some of the key concepts described above by using the Kerr metric as an example. As described in Section 2.4.3, the exterior of the Kerr spacetime (i.e., outside the outer event horizon) is usually expressed in terms of Boyer–Lindquist coordinates $\{t, r, \theta, \phi\}$. One may instead express the metric in *canonical coordinates* $\{\tau = t - a\phi, r, y = a \cos \theta, \psi = \frac{\phi}{a}\}$, where a is the spin of the black hole. The coordinates r and y are related to the eigenvalues of the principal tensor, and the coordinates τ and ψ are related to the Killing vectors which correspond to stationarity and axisymmetry, respectively. In canonical coordinates, the line element reads

$$ds^2 = g_{ab} dx^a dx^b = -\frac{\Delta_r}{\Sigma} (d\tau + y^2 d\psi)^2 + \frac{\Delta_y}{\Sigma} (d\tau - r^2 d\psi)^2 + \frac{\Sigma}{\Delta_r} dr^2 + \frac{\Sigma}{\Delta_y} dy^2, \quad (2.153)$$

where $\Sigma = r^2 + y^2$, $\Delta_r = r^2 - 2Mr + a^2$ and $\Delta_y = a^2 - y^2$. Replacing the canonical coordinates by the Boyer–Lindquist coordinates, we arrive at the familiar form of the Kerr metric (2.86).

The Kerr metric admits a pair of Killing vectors, namely $\xi_{(\tau)} = \partial_\tau = \partial_t$ and $\xi_{(\psi)} = \partial_\psi = a^2 \partial_t + a \partial_\phi$. As discussed above, the Kerr spacetime admits a principal tensor (i.e.,

a closed conformal Killing–Yano two-form) which satisfies $h_{ab;c} = g_{ca}\xi_{(\tau)b} - g_{cb}\xi_{(\tau)a}$. The primary Killing vector can be derived from the principal tensor by noting that $\xi_{(\tau)}^a = \frac{1}{3}h^{ba}{}_{;b}$. In canonical coordinates, the principal tensor takes the form

$$h = y \, dy \wedge (d\tau - r^2 d\psi) - r \, dr \wedge (d\tau + y^2 d\psi). \quad (2.154)$$

The fact that h is closed means that it can be expressed in terms of a one-form potential b ($h = db$) which in canonical coordinates takes the form

$$b = -\frac{1}{2} [(r^2 - y^2) d\tau + r^2 y^2 d\psi]. \quad (2.155)$$

The Hodge dual of the principal tensor is the Killing–Yano tensor $f = *h$; in component form this is $f_{ab} = \frac{1}{2}\epsilon_{abcd}h^{cd}$, where ϵ_{abcd} is the Levi-Civita tensor. In canonical coordinates, the Killing–Yano tensor is

$$f = r \, dr \wedge (d\tau - r^2 d\psi) + y \, dy \wedge (d\tau + y^2 d\psi). \quad (2.156)$$

The Killing–Yano tensor has a totally antisymmetric derivative: $f_{ab;c} = f_{[ab;c]}$. The “square” of the Killing–Yano tensor is the Killing tensor $K_{ab} = f_{ac}f_b{}^c$, which satisfies the Killing tensor equation $K_{(ab;c)} = 0$. The “square” of the conformal Killing–Yano tensor is the conformal Killing tensor $Q_{ab} = h_{ac}h_b{}^c$, which satisfies the conformal Killing tensor equation $Q_{(ab;c)} = g_{(ab}q_{c)}$, where the rank-one tensor q_a can be expressed in terms of the principal tensor and the primary Killing vector as $q_a = h_{ab}\xi_{(\tau)}^b$. The secondary Killing vector is related to the primary Killing vector and the Killing tensor via $\xi_{(\psi)} = -K^a{}_b \xi_{(\tau)}^b$.

The conserved quantities on phase space which are induced by the Killing vectors $\xi_{(\tau)}^a$ and $\xi_{(\psi)}^a$, the Killing tensor K_{ab} , and the metric tensor g_{ab} (which is itself a trivial Killing tensor) are all independent and in involution (i.e., they mutually Poisson commute). This fact ensures that the geodesic motion on Kerr spacetime is completely integrable (see Section 2.6.3). The separability of geodesic motion on Kerr spacetime was first demonstrated by Carter [47]. Subsequently, in 1973, Floyd [48] and Penrose [49] demonstrated that the Kerr spacetime admits a rank-two Killing tensor K^{ab} which can be constructed from the dual of the closed conformal Killing–Yano tensor, or more directly from the Killing–Yano tensor. The Killing tensor is responsible for the existence of a fourth “hidden” constant of the motion, which can be expressed in the form $K = K^{ab}p_a p_b$; this is referred to as the *Carter constant*. Thus, the principal tensor reveals “hidden” symmetries of Kerr spacetime which guarantee that the geodesic dynamics are integrable in the sense of Liouville.

There are a number of further important consequences for geodesic motion on Kerr spacetime, which are related to the existence of the principal tensor f_{ab} . Firstly, the vector $L^a = f^{ab}p_b$ satisfies the equation of parallel transport $p^b \nabla_b L^a = 0$, so gyroscopic precession is straightforward [52]; this fact will be important in Chapter 7. Secondly, gravitational Faraday rotation of the polarisation plane of light is trivial: the complex quantity $z = (f^{ab} + ih^{ab})p_a v_b$ is constant along a null ray with momentum p_a for any parallel-transported vector v^a [49].

The hidden symmetry brought about by the principal tensor is also implicated in a number of other important results relating to the separability of wave equations on

Kerr spacetime. Firstly, the Klein–Gordon equation $\square\Phi = g^{ab}\nabla_a\nabla_b\Phi = 0$ for a scalar field Φ admits separable solutions of the form $\Phi = e^{-i\omega t}e^{im\phi}R(r)S(\theta)$ in Boyer–Lindquist coordinates, where $R(r)$ and $S(\theta)$ obey second-order ordinary differential equations [53]. Secondly, the solutions to the Dirac equation are separable, and the governing equations reduce to a set of two coupled ordinary differential equations [54, 55]. Thirdly, Maxwell’s equations admit separable solutions for the Maxwell scalars of extreme spin weight ± 1 , given by $\phi_0 = F_{ab}k^am^b$ and $\phi_2 = F_{ab}n^a\bar{m}^b$, where $\{k, n, m, \bar{m}\}$ is a certain complex null tetrad (see Section 2.5.1), and F_{ab} is the Faraday two-form (see Sections 2.3.2 and 3.3.1) [58]. Finally, the equations governing gravitational perturbations admit separable solutions for the Weyl scalars of extreme spin weight ± 2 , given by $\Psi_0 = C_{abcd}k^am^bk^cm^d$ and $\Psi_4 = C_{abcd}\bar{m}^an^b\bar{m}^cn^d$ (see Section 2.5.4) [59].

2.6.5 Non-integrable systems and chaos

In Section 2.6.3, we saw that, in Hamiltonian systems, the existence of a maximal number of independent conserved quantities which are in involution guarantees the Liouville integrability of motion in phase space. In such circumstances, the equations of motion can be solved by quadrature, either analytically or using numerical techniques. In general, however, this is almost never possible.

When a Hamiltonian system does not admit a maximal set of independent Poisson-commuting integrals of motion, i.e., the system is non-integrable, the trajectories have much less restriction on their motion in phase space: they are not required to lie on invariant tori. In these situations, which are generic, one observes *sensitive dependence on initial conditions* in (regions of) phase space. Inevitably, *chaotic* behaviour will occur in non-integrable systems.

In this section, we review some fundamental aspects of deterministic chaos, with a particular emphasis on chaotic scattering and open Hamiltonian systems; we also discuss some of the hallmarks of chaotic motion, such as fractal structures in phase space.

Deterministic chaos and fractals in dynamical systems

The solutions of simple non-linear dynamical systems can behave in an extremely complicated fashion [95, 96]. Heuristically, a chaotic solution is aperiodic, and initially neighbouring solutions diverge exponentially in time. The former property indicates that the solution is *irregular* or *disordered* in some sense, and the latter property, known as *sensitive dependence on initial conditions*, suggests that even the smallest changes in the initial state of the system can yield wildly different outcomes. Sensitive dependence on initial conditions makes long-term predictions of the system very difficult, despite the deterministic nature of equations which govern the dynamics. This type of behaviour is known as *deterministic chaos* (or simply *chaos*). A precise mathematical definition of chaos in the context of maps (i.e., discrete-time dynamical systems) is given by Devaney [95]. Often, the term “chaos” is used to refer only to motion on compact manifolds. Compact means that systems with escapes are excluded; in such cases, the term “chaotic scattering” is used, rather than “chaos”. Extensive reviews of deterministic chaos can be found in a

range of textbooks and articles, e.g. [95–100].

Chaotic behaviour is typically characterised using a variety of mathematical methods, which can be either qualitative or quantitative in nature. For example, a common diagnostic of chaos is the calculation of *Lyapunov exponents* [97], which quantify the rate of separation of initially neighbouring trajectories.

Another hallmark of chaos is the presence of fractal structures in phase space [112]. Loosely speaking, a fractal is a complicated geometrical object which exhibits irregular structure at arbitrarily small scales [97, 113, 114]. Much like with the notion of chaos, there have been a variety of attempts to give a precise mathematical definition of a fractal; however, many of these definitions are not general enough to capture the rich structure of many sets which are intuitively fractals. Falconer [113, 114] considers a fractal to be a set F in Euclidean space which exhibits all or most of the following characteristic features.

- (i) The set F has a fine structure, i.e., it exhibits irregular detail at arbitrarily small scales.
- (ii) Due to its “irregularity”, F cannot be described using the language of calculus or traditional Euclidean geometry, either locally or globally.
- (iii) In many cases, F will exhibit self-similarity or self-affinity, perhaps in a statistical or approximate manner.
- (iv) The fractal dimension of F is strictly greater than its topological dimension (for some suitable definition of fractal dimension, e.g. the Hausdorff dimension).
- (v) A simple, perhaps recursive, procedure (e.g. an iterated function system) can be used to generate F .
- (vi) The set F will often have a “natural” appearance.

There exist a multitude of fractals which possess some or all of the properties listed above. A particularly interesting class of fractals are those which are invariant under simple families of transformations, including self-similar, self-affine, approximately self-similar, and statistically self-similar fractals. Some of the best-known examples of fractals are self-similar, e.g. the middle-thirds Cantor set, the von Koch curve, the Sierpinski triangle and the Sierpinski carpet. Further popular examples of fractals are those which occur as the attractors of (mathematical) dynamical systems, such as the Mandelbrot set or Julia sets, both of which are the result of complex iterated function systems.

Fractals, however, are not only a mathematical curiosity: they arise naturally in non-linear dynamical systems, in such a way that the two concepts are intimately related. For a detailed account of the relationship between fractal structures and non-linear dynamics, we refer the reader to the review by Aguirre *et al.* [112].

Poincaré sections

We now turn to an important technique in chaotic dynamics (in the case of bounded motion), originally suggested by Poincaré [115], which permits us to study the breakdown in

integrability, the destruction of invariant tori, and the nature of motion in the gaps between the destroyed tori [103]. This method is particularly well suited to Hamiltonian systems in two dimensions (although it can be employed in higher-dimensional systems). In such situations, the phase space is spanned by four canonical coordinates, say $\{q^1, q^2, p_1, p_2\}$, and the motion is confined to an *energy hypersurface* \mathcal{E} , a three-dimensional hypersurface described by the Hamiltonian constraint $H(q^1, q^2, p_1, p_2) = h$, where h is the (conserved) total energy. (We note that, in the case of null geodesic motion in general relativity, the Hamiltonian constraint is $H = 0$ for *all* trajectories in phase space.)

Rather than trace out the orbits which wander through the three-dimensional energy hypersurface, it is advantageous to study a two-dimensional slice through the surface; this is known as a *Poincaré section* (or *surface of section*). A typical choice for the section is the surface \mathcal{S}_1 , given by the intersection of the hypersurface \mathcal{E} with $q^2 = 0$; this surface of section has coordinates (q^1, p_1) . (Of course, one may define the surface of section \mathcal{S}_2 analogously, or choose a more general two-dimensional surface which intersects \mathcal{E} .) Specifying a point on \mathcal{S}_1 then completely determines a dynamical trajectory (modulo a direction): q^1, p_1 and $q^2 = 0$ are all specified, and a rearrangement of the Hamiltonian constraint $H(q^1, q^2, p_1, p_2) = h$ allows us to determine p_2 up to a sign (since H is assumed to be quadratic in the momenta). By convention, this sign is taken to be positive.

A trajectory which begins on \mathcal{S}_1 will subsequently cross it repeatedly as it moves on \mathcal{E} . The *Poincaré map* $P: \mathcal{S}_1 \rightarrow \mathcal{S}_1$ maps a trajectory that begins on \mathcal{S}_1 to its next point of intersection with \mathcal{S}_1 . The Poincaré map reduces a continuous-time dynamical system (a flow) on a three-dimensional energy surface to a discrete-time dynamical system (a map) on a two-dimensional surface of section. An important property of the map P is that it is *area preserving* [103].

The intersections of an orbit with the Poincaré section reveal whether the motion is integrable. If this is the case, the system will not explore the whole of the three-dimensional hypersurface \mathcal{E} ergodically; rather, the motion will be restricted to a two-torus (see Section 2.6.3), which intersects the surface of section in a closed curve. All iterates of this orbit under the action of the Poincaré map must lie on this curve. For *periodic* orbits, some iterate will coincide with the initial position on the surface of section, with the order of the iterate depending on the commensurability of the system's characteristic frequencies; such orbits are fixed points of the map P^k , where k is the order of the iterate. Most curves, however, are formed by the intersection of irrational (non-resonant) tori with the Poincaré section. All of these curves are *invariant curves* of the Poincaré map, as they are mapped to themselves by P .

Consider a perturbed system with Hamiltonian $H = H_0 + \varepsilon H_1$, where H_0 is the Hamiltonian of an integrable system. When $\varepsilon \ll 1$, the system is said to be *near-integrable*. As one increases ε , the system undergoes a transition from integrability to non-integrability, which obeys the Kolmogorov–Arnold–Moser (KAM) theorem [116]. This says that for sufficiently small perturbations ε , almost all non-resonant tori of the integrable system survive, but are slightly deformed. Such tori are known as *KAM tori*. Here, “almost all” means that the measure of the complement of the KAM tori is small, and tends to zero with the magnitude of the perturbation ε . The surviving tori are those which are

“sufficiently irrational” [101, 102].

In the case of non-integrable systems, the invariant tori do not exist, and orbits explore a three-dimensional region of the energy hypersurface \mathcal{E} . The intersection of such orbits with the Poincaré section do not lie on a curve, but instead cover a two-dimensional region of \mathcal{S}_1 .

A typical Poincaré section will exhibit invariant curves, resonant “islands” (i.e., intersections of the nested tori with the surface of section), and chaotic regions; the latter may border the islands, or fill more extensive regions of the surface of section. The ordered phase space of the integrable system becomes highly intricate under perturbations. As the perturbation is increased, the system moves further away from integrability; more and more invariant tori are destroyed, and the chaotic regions grow and merge with one another, filling up the surface of section.

The usefulness of Poincaré sections is best demonstrated by considering an example. In Appendix C, we illustrate the method by applying it to the Hénon–Heiles Hamiltonian system [117] as a pedagogical example. This will be of particular importance in Chapter 6, in which we use Poincaré sections to analyse stable bounded null geodesics around a pair of black holes.

Chaotic scattering

We now turn our attention to the case of *unbounded* chaotic motion, in which trajectories are permitted to reach infinity. Motion of this type is referred to as *chaotic scattering*. In its most general sense, a *scattering system* is the problem of obtaining a relationship between an “input variable” (e.g. the impact parameter, or the initial angle), which characterises an *initial state* of a dynamical system (usually a Hamiltonian system), and an “output variable” (e.g. the scattering angle, or some discrete label), which characterises the *final state* of the dynamical system [118]. Typically, a scattering problem deals with the motion of a particle described by the Hamiltonian

$$H(q, p) = H_0(q, p) + V(q), \quad (2.157)$$

where V is a potential which is zero (or sufficiently small) outside of some bounded region of space which is referred to as the *scattering region* [99]. The particle can effectively be treated as a free particle, governed by the Hamiltonian H_0 , far from the scattering region.

Chaotic scattering is the study of scattering systems which exhibit extreme sensitive dependence on initial conditions. Eckhardt [119] defines scattering in a Hamiltonian system as *irregular* (or *chaotic*) if there exists, on some manifold of initial data, an infinity of distinct scattering singularities of zero (Lebesgue) measure, which are typically arranged into a fractal set. In this context, a *scattering singularity* is an initial value (input) for which the scattering process (output) is undefined. In regions of chaotic scattering (i.e., in neighbourhoods of scattering singularities), small perturbations in the initial conditions lead to completely different outcomes (final states) in the scattering process.

In a general scattering system, the *classical repellor*, denoted Ω_R , is the set of all scattering singularities, which correspond to non-escaping orbits [100, 120]. In general, Ω_R may be a simple set which contains a finite or countable number of scattering singularities,

corresponding to periodic orbits. On the other hand, Ω_R may be a fractal set (of measure zero) which comprises a countable set of periodic orbits, and an uncountable set of aperiodic orbits. The former case corresponds to regular scattering, whilst the latter represents chaotic scattering. If the repeller is a fractal, it is referred to as a *strange repeller*.

Chaotic scattering occurs in wide variety of contexts, including the three-disc model [119, 120], and the Hénon–Heiles Hamiltonian system of celestial mechanics [117]. The phenomenology of chaotic scattering is reviewed in a range of references, e.g. [100, 121, 122].

Open Hamiltonian systems and exit basins

An *open Hamiltonian system* is a Hamiltonian dynamical system in which the particle is permitted to escape from the scattering region to infinity, via one or more *escapes* (or *exits*) in phase space [112].

The dynamics of open Hamiltonian systems can be understood in more detail by constructing exit basins in phase space, in a fashion analogous to the basins of attraction for dissipative dynamical systems [123–125]. Dissipative dynamical systems often possess more than one *attractor*. A *basin of attraction* is defined as a subset of state space such that all initial conditions chosen from this set asymptote towards a given attractor in forward time. In an open Hamiltonian system with $N \geq 2$ exits in phase space, the *exit basin* corresponding to the i th exit in phase space is defined to be the set of all initial conditions which escape the scattering region through exit i ; see Section 2.6.5 and references therein.

When more than one basin is present, the basins will be separated by a *basin boundary*. A point p is a *boundary point* of the basin B if $b_\varepsilon(p) \cap B \neq \emptyset$ and $b_\varepsilon(p) \cap B^c \neq \emptyset$ for all $\varepsilon > 0$, where B^c is the complement of the set B , and $b_\varepsilon(p)$ denotes the open ball of radius ε centred on p . The basin boundary, denoted ∂B , is then the set of all boundary points of B .

The exit basins of an open Hamiltonian system may be visualised by integrating the equations of motion for a fine grid of initial conditions, then colour-coding the initial data according to the escape through which the corresponding trajectories leave the scattering region. The result of this procedure is a so-called *exit basin diagram* [112], which provides a useful tool with which one can analyse the global dynamics of open Hamiltonian systems. In Hamiltonian systems with more than two degrees of freedom (i.e., the phase space has dimension four or higher), one must use a Poincaré section to realise the exit basin diagram. This involves choosing an one- or two-dimensional surface in phase space, from which the initial conditions are chosen. The trajectories are then integrated forwards in time, and the initial data on the surface are colour-coded according to the final state of the corresponding trajectory. (Note that, in this case, the motion is unbounded, so the Poincaré map is not well-defined.)

If the system under consideration is chaotic, one typically observes *fractal* structures in the exit basin diagrams; in particular, the boundaries between any two exit basins may be fractal. In this case, there exist a variety of methods which can be used to characterise (both qualitatively and quantitatively) the nature of these fractal structures and the chaotic dynamics; see the review by Aguirre *et al.* [112] for further details.

Chaos in general relativity

General relativity is an inherently non-linear theory. The Einstein field equations (2.75), which govern the geometry of spacetime, are a set of coupled non-linear partial differential equations; this means that the spacetime dynamics can, in principle, exhibit chaotic behaviour. This has been the subject of a range of studies over the last four decades, with a particular emphasis on the study of chaos in the spacetime dynamics of cosmological Bianchi IX (Mixmaster) models [126–128].

During the last thirty years, investigations of chaos in general relativity have followed an alternative route: the analysis of geodesic motion on a *fixed* spacetime background. It is common to start with a background on which the geodesic motion is integrable (e.g. Schwarzschild or Kerr), and then consider the effect of various types of “perturbations”. These perturbations typically reduce the symmetry required for complete integrability, and rich chaotic behaviour emerges naturally. Such behaviour can be characterised using a variety of methods. The fact that geodesic motion in general relativity is amenable to Hamiltonian methods means that existing techniques can be borrowed from the theory of non-linear dynamical systems. However, care must be taken to ensure that these descriptors of chaos are coordinate-invariant. Here, we review some key studies of chaos in geodesic motion in general relativity. (Further details can be found in the comprehensive summary presented by Semerák and Suková in [129].)

A range of authors have considered the effect of endowing a test particle with spin, looking at the dynamics of the Mathisson–Papapetrou–Dixon equations on Schwarzschild spacetime [130], and Kerr spacetime [131, 132]. Others have focussed on chaos in gravitational waves. In 1992, Bombelli and Calzetta [133] studied the influence of small periodic perturbations (which serve as a model for weak gravitational waves) on Schwarzschild spacetime. Further studies of chaos in gravitational waves were carried out by Levin [134] and others [135, 136].

Various authors have looked at perturbing the background itself by endowing it with higher multipoles. For example, Vieira and Letelier [137] analysed the effect of perturbations by quadrupolar and octupolar moments on the timelike geodesics of Schwarzschild spacetime. Later, by considering exit basins in phase space, de Moura and Letelier [138] considered the escape dynamics of timelike and null geodesics from a static black hole endowed with dipolar, quadrupolar and octupolar moments.

Rather than artificially superposing a regular background with higher order multipoles, one may also analyse exact solutions which incorporate these multipoles. Such solutions are typically stationary and axisymmetric. In 1996, Sota *et al.* [139] studied the Zipoy–Voorhees family of solutions, which includes Minkowski, Schwarzschild and Chazy–Curzon spacetimes as special cases. The authors used techniques such as Poincaré sections and Lyapunov exponents to diagnose chaos in the motion of free test particles.

Motivated by astrophysical scenarios, a range of studies have looked at relativistic core–shell models, such as the superposition of a black hole and a self-gravitating disc or ring [138]; see [129] and references therein for a review.

Finally, there has been much interest in the analysis of regular and chaotic dynamics in

the motion of free particles around two (or more) fixed centres. This began with the work of Contopoulos [140–142], who demonstrated that the motion of free particles around two fixed black holes in the Majumdar–Papapetrou spacetime is chaotic. This topic will be of central importance in Chapters 4–6 of this thesis; a review of key contributions to this field will be provided there.

Chapter 3

Gravitational lensing

3.1 Introduction

Gravitational lensing is a general term which accounts for all effects of a gravitational field on the propagation of electromagnetic radiation [7]. Typically, this electromagnetic radiation is described by light rays, using a high-frequency (geometrical optics) approximation. The gravitational field which affects these light rays can itself be described using a variety of approaches.

In situations where the gravitational field is “weak”, quasi-Newtonian approximations may be employed to reduce the formalism of general relativity to a Newtonian setting; this is often referred to as *perturbative lensing*. The review by Wambsganss [70] provides a comprehensive overview of observational aspects and key methods of perturbative lensing, with an emphasis on gravitational lensing as an astrophysical tool. Due to the fact that the quasi-Newtonian formalism applies in situations where the magnitude of the Newtonian gravitational potential is small (i.e., in the weak-field limit), perturbative lensing is sometimes referred to as *weak-field gravitational lensing*.

On the other hand, one may analyse *lensing from a spacetime perspective* (also known as *non-perturbative lensing*). In this approach, the propagation of light rays is described by null geodesics on a (four-dimensional) Lorentzian spacetime, which is a solution to the Einstein field equations of general relativity. Of course, the latter condition may be dropped in favour of an alternative theory of gravity, if so desired. For an in-depth account of gravitational lensing from a spacetime perspective, see the review by Perlick [7], and references therein. Non-perturbative lensing can be used to describe gravitational phenomena in both weak- and strong-field regimes. In situations where we are only concerned with the latter, we use the terminology *strong-field gravitational lensing*.

In this chapter, we present a brief review of gravitational lensing. Section 3.2 concerns perturbative lensing. We begin with a short historical account of the origins of the theory, illustrated by important theoretical and observational milestones; we then review the list of observed gravitational lensing phenomena; and we present a brief account of the mathematical formalism of perturbative lensing. In Section 3.3, we discuss key elements of the theory of non-perturbative gravitational lensing, within the framework of four-dimensional general relativity. First, we review Maxwell’s theory of electromagnetism in

curved spacetime. We then describe the (leading-order) geometrical optics approximation, which allows the electromagnetic field to be described in terms of null rays on curved spacetime. We conclude with an introduction to some key features of gravitational lensing by black holes, including light-rings (i.e., unstable photon orbits) and black hole shadows.

3.2 Perturbative lensing

3.2.1 History of gravitational lensing

The idea of light rays being affected by a gravitational field can be traced back to Isaac Newton. In his 1704 treatise *Opticks*, Newton queries whether (gravitational) bodies act upon light at a distance, and by their action bend its rays [143]. It was not until 1784 that Henry Cavendish calculated the deflection angle of light due to a point mass, using Newton’s gravitational theory and Newtonian optics, in which light is assumed to be composed of corpuscles that are affected by a gravitational field in the same way as material particles.

Interestingly, it was Cavendish’s correspondence with his friend and colleague John Michell which sparked the former’s interest in the effect of gravity on light. In 1783, Michell wrote a letter to Cavendish [144] which proposed a method to measure the mass of stars by measuring the reduction in the speed of the light corpuscles which propagate from the star to the Earth. It was this paper and the correspondence with Michell that prompted Cavendish to consider the deflection of light due to a gravitational field. Cavendish, however, never published his calculation of light deflection; evidence of this work was only found after Cavendish’s death [145].

In his 1783 letter, Michell also proposed that there might exist “dark stars”, objects which are so dense that not even light can escape from their surface; such a body would therefore be invisible, and only detectable by its gravitational influence on nearby stars [146]. This proposal is now widely quoted as the origin of the concept of a black hole – an object so dense that not even light can escape its gravitational pull. Over a decade later, Pierre-Simon Laplace, independently of Michell, calculated the conditions required for the escape velocity from the surface of a body to be greater than the velocity of light [147]. Curiously, it is the work of Laplace that inspired Johann Georg von Soldner to provide a detailed calculation of the deflection angle suffered by a light ray due to the gravitational attraction of a massive body [148, 149]: von Soldner found that a ray passing close to a body of mass M is bent through the angle

$$\alpha_{\text{N}} = \frac{2GM}{c^2 b}, \quad (3.1)$$

where b is the ray’s perpendicular impact parameter, G is Newton’s gravitational constant, and c is the speed of light. For a ray which grazes the limb of the Sun, the deflection angle is $\alpha_{\odot} \approx 0.87$ arcsec. An account of the contributions of Cavendish and von Soldner to the field of gravitational lensing is provided by Will [63].

In 1911, Albert Einstein calculated the deflection angle of light by the Sun using the equivalence principle of his special theory of relativity [150]. The result of Einstein’s

calculation matched that of von Soldner, of whom he knew nothing about. A few years later, in 1914, at the request of Einstein, the German astronomer Erwin Freundlich led an expedition to the Crimean peninsula, with the aim of observing light deflection by the Sun during a total solar eclipse. Unfortunately, the outbreak of the First World War prevented Freundlich’s team from completing the expedition [71].

Equipped with his recently developed general theory of relativity, Einstein revisited the calculation of light deflection in 1915. Using a linearised version of the field equations, Einstein found that the deflection angle was actually twice the Newtonian prediction of Cavendish and von Soldner [151, 152]:

$$\alpha_E = \frac{4GM}{c^2 b}, \quad (3.2)$$

with $\alpha_\odot \approx 1.75$ arcsec for a ray which grazes the Sun’s surface. Using the Schwarzschild solution to Einstein’s field equations, which describes a spherically symmetric non-rotating source (see Section 2.4), the general-relativistic deflection angle can be calculated exactly in terms of an elliptic integral [153, 154].

The theoretical prediction of light deflection by the Sun due to general relativity was verified experimentally during a total solar eclipse in May 1919, by a team of astronomers led by Sir Frank Watson Dyson and Sir Arthur Eddington [155]. The accuracy of the 1919 observations was sufficient to distinguish between the predictions of Newtonian gravity and Einstein’s theory of general relativity. This observation helped to establish the general relativity as a fundamental physical theory of gravitation.

In 2004, the deflection by the Sun of radio waves emitted by distant compact radio sources was measured using very-long-baseline interferometry; this confirmed the general-relativistic prediction of 1.75 arcsec to within 0.2% [156]. More recently, during the total solar eclipse of August 2017, Einstein’s value was observed for visible light, with an accuracy of 3% [157].

3.2.2 Gravitational lensing phenomena

The success of the 1919 observations of light deflection by the Sun marked the birth of gravitational lensing. In the decades that followed, gravitational lensing was developed as a theoretical tool.

Chwolson [64] proposed the idea of a “fictitious double star” in 1924, and described the symmetric case in which the source, lens and observer are perfectly aligned, resulting in a ring-like image of the source. In 1936, Einstein described the result of the latter lens configuration as a “luminous circle” [65]; today this phenomenon is known as an *Einstein ring*. Curiously, Einstein disregarded his results, claiming that there was no hope of observing such an effect directly.

One year later, influenced by the work of Einstein, Zwicky [158] posited that distant galaxies could act as sources and as gravitational lenses. Zwicky noted that galaxy–galaxy lensing could be used as an observational tool to test general relativity, and that these lensing effects could be used as a natural telescope to observe distant galaxies, due to the magnification effect of gravitational lensing.

The field of gravitational lensing remained dormant until the early 1960s, when interest in the subject was reignited by the theoretical developments of Klimov [67], Liebes [68] and Refsdal [66]. In particular, Refsdal established the quasi-Newtonian formalism for gravitational lensing, based on the assumptions of weak gravitational fields and small deflection angles; and demonstrated that gravitational lensing could be used as a tool to measure the Hubble constant by measuring the time delay between two images of a single source [159]. Interestingly, earlier that decade, Sachs [160] provided a fully covariant formulation of how a bundle of light rays is distorted by a gravitational field in the context of gravitational radiation; this will be discussed in Section 3.3.

The first detection of a quasar – a luminous, compact and very distant source – came in 1963 [161]. Such an object would be an ideal source for Zwicky’s galaxy–galaxy lensing, and it was hoped that the observation of gravitational lensing would soon follow. Fifteen years later, Walsh *et al.* [69] discovered the “twin quasar” Q0957+561, and confirmed that this was in fact one quasar, of which two images are produced by gravitational lensing. This was the first detection of a gravitational lensing effect since the observation of light deflection by the Sun in 1919, and signalled the birth of gravitational lensing as an observational field [70].

To date, the list of observed lensing phenomena includes multiply imaged quasars; rings; giant luminous arcs; image distortion by weak (or statistical) lensing; and galactic microlensing. Detailed reviews of these lensing effects can be found in a range of sources; see for example [70, 71, 162, 163] and references therein.

3.2.3 Perturbative lensing formalism

In this section, we review some basic aspects of the mathematical formalism of perturbative gravitational lensing. In its most general sense, solving for the motion of light rays on an arbitrary curved spacetime is a non-trivial theoretical task. However, in many cases, one may assume that the spacetime geometry is well-described by the Friedmann–Lemaître–Robertson–Walker metric, and that the matter inhomogeneities which are responsible for gravitational lensing can be treated as (small) local perturbations. The mathematical formalism of linearised gravity and perturbative gravitational lensing is reviewed in a range of sources, e.g. [70, 71, 162–164]. Here, we follow the discussion presented by Carroll [165].

For simplicity, we restrict our attention to cases where the background spacetime is well described by the flat Minkowski metric $ds^2 = \eta_{ab} dx^a dx^b = -dt^2 + dx^2 + dy^2 + dz^2$. The fact that the gravitational field is weak means that one may express the spacetime metric as

$$g_{ab} = \eta_{ab} + h_{ab}, \quad |h_{ab}| \ll 1, \quad (3.3)$$

where h_{ab} is regarded as a small perturbation, and $|h_{ab}|$ denotes the magnitude of a typical non-zero component of the perturbation. (Note that we returned to units in which $c = 1 = G$ for simplicity.) In linearised gravity, we keep only terms which are linear in the perturbation h_{ab} ; higher-order terms are discarded. It follows that $g^{ab} = \eta^{ab} - h^{ab}$, where $h^{ab} = \eta^{ac} \eta^{bd} h_{cd}$. A key consequence of this is that spacetime indices may be lowered and raised using the Minkowski metric and its inverse, as any corrections will be higher than

first order in the perturbation h_{ab} .

Typically, one assumes that the source of the gravitational field is static and Newtonian. Moreover, the gravitational source is modelled as dust, a perfect fluid with vanishing pressure. In this approximation, the perturbed metric reads

$$ds^2 = g_{ab}dx^a dx^b = -(1 + 2\Phi) dt^2 + (1 - 2\Phi) (dx^2 + dy^2 + dz^2), \quad (3.4)$$

in Cartesian coordinates $\{t, x, y, z\}$. Comparing (3.3) and (3.4), we see that $h_{ab} = -2\Phi \delta_{ab}$ and $|\Phi| \ll 1$. The Newtonian gravitational potential Φ satisfies Poisson's equation $\nabla^2 \Phi = 4\pi\rho$, where ρ is the mass density of the gravitational source in the rest frame of the dust.

We are interested in the propagation of light rays on the geometry (3.4). We therefore wish to solve the perturbed null geodesic equation for a ray $x^a(\lambda)$, where λ is an affine parameter. As in the case of the metric tensor (3.3), we decompose the ray into a background path (in Minkowski spacetime) plus a perturbation:

$$x^a(\lambda) = x_{(0)}^a(\lambda) + x_{(1)}^a(\lambda). \quad (3.5)$$

Here, $x_{(0)}^a(\lambda)$ solves the geodesic equation on Minkowski spacetime; the path is therefore a straight line. To solve for the perturbation $x_{(1)}^a(\lambda)$, we adopt the *Born approximation* [71], i.e., we evaluate all quantities along the background path rather than the actual path. This procedure is valid because we assume that $|\Phi| \ll 1$, so we expect the deviation to be small.

We introduce the notation

$$u^a = (u^0, \mathbf{u}) = \dot{x}_{(0)}^a, \quad v^a = (v^0, \mathbf{v}) = \dot{x}_{(1)}^a, \quad (3.6)$$

where an overdot denotes differentiation with respect to the affine parameter λ . The path (3.5) must be null, i.e., $g_{ab}\dot{x}^a\dot{x}^b = 0$. This equation must be solved order-by-order in the perturbation. At zeroth order, this is nothing more than $\eta_{ab}u^a u^b = 0$, which says that u^a a null vector with respect to the background (Minkowski) spacetime; this condition may be written

$$-(u^0)^2 + |\mathbf{u}|^2 = 0. \quad (3.7)$$

We define $u^2 = (u^0)^2 = |\mathbf{u}|^2$. At first order in the perturbation, the null condition is $2\eta_{ab}u^a v^b + h_{ab}u^a u^b = 0$, which may be written in the form

$$-u^0 v^0 + \mathbf{u} \cdot \mathbf{v} = 2u^2 \Phi. \quad (3.8)$$

The perturbed geodesic equation is

$$\ddot{x}^a + \Gamma^a_{bc}\dot{x}^b\dot{x}^c = 0, \quad (3.9)$$

where, to first order in the perturbation, the Christoffel symbols are given by

$$\Gamma^a_{bc} = \frac{1}{2}g^{ad}(\partial_b g_{dc} + \partial_c g_{bd} - \partial_d g_{bc}) = \frac{1}{2}\eta^{ad}(\partial_b h_{dc} + \partial_c h_{bd} - \partial_d h_{bc}). \quad (3.10)$$

For the static Newtonian gravitational field with line element (3.4), the only non-zero Christoffel symbols are

$$\Gamma^0_{i0} = \partial_i \Phi, \quad \Gamma^i_{00} = \partial_i \Phi, \quad \Gamma^i_{jk} = \delta_{jk} \partial_i \Phi - \delta_{ik} \partial_j \Phi - \delta_{ij} \partial_k \Phi. \quad (3.11)$$

At zeroth order, (3.9) reduces to $\dot{u}^a = 0$. This tells us that $x_{(0)}^a(\lambda)$ is a straight line (i.e., a geodesic in Minkowski spacetime), as anticipated. At first order in the perturbation, (3.9) yields $\dot{v}^a + \Gamma^a_{bc} u^b u^c = 0$. The temporal and spatial components of this equation are

$$\dot{v}^0 = -2u(\mathbf{u} \cdot \nabla\Phi), \quad \dot{\mathbf{v}} = -2u^2 \nabla_{\perp}\Phi, \quad (3.12)$$

where the *transverse gradient* is defined as

$$\nabla_{\perp}\Phi = \nabla\Phi - \nabla_{\parallel}\Phi = \nabla\Phi - \frac{1}{u^2}(\mathbf{u} \cdot \nabla\Phi)\mathbf{u}. \quad (3.13)$$

To first order in the perturbation, \mathbf{u} is orthogonal to \mathbf{v} . To demonstrate this, one may obtain an expression for v^0 by integrating the first equation of (3.12) along the background path:

$$v^0 = -2u \int (\mathbf{u} \cdot \nabla\Phi) d\lambda = -2u \int (\nabla\Phi \cdot \dot{\mathbf{x}}_{(0)}) d\lambda = -2u \int \nabla\Phi \cdot d\mathbf{x}_{(0)} = -2u\Phi, \quad (3.14)$$

where the integration constant is chosen to ensure that $v^0 = 0$ when $\Phi = 0$, which comes from the fact that the deviation vanishes in the absence of gravitational sources. Inserting (3.14) into (3.8) shows that $\mathbf{u} \cdot \mathbf{v} = u^0 v^0 + 2u^2 \Phi = 0$, as claimed. Alternatively, one may note that

$$\frac{d}{d\lambda}(\mathbf{u} \cdot \mathbf{v}) = \dot{\mathbf{u}} \cdot \mathbf{v} + \mathbf{u} \cdot \dot{\mathbf{v}} = -2u^2(\mathbf{u} \cdot \nabla_{\perp}\Phi) = 0, \quad (3.15)$$

where we have used (3.12). The quantity $\mathbf{u} \cdot \mathbf{v}$ is therefore conserved along the background ray. Of course, there is some freedom in the choice of initial direction for the perturbation. Choosing this to point along the background ray is equivalent to rescaling the affine parameter λ . Our choice is to take $\mathbf{u} \cdot \mathbf{v} = 0$ initially; (3.15) ensures that this remains so along the background ray.

The deflection angle $\boldsymbol{\alpha}$ (discussed in Section 3.2.1) is the amount by which the spatial wave vector $\dot{\mathbf{x}}$ is deflected as it travels from the light source to the observer in the gravitational field (3.4). The deflection angle is a two-dimensional vector which lies in the plane perpendicular to \mathbf{u} ; this may be expressed in the form

$$\boldsymbol{\alpha} = -\frac{\Delta\mathbf{v}}{u}, \quad (3.16)$$

where the minus sign is chosen to account for the fact that the deflection is measured by an observer looking backwards along the null ray. The quantity $\Delta\mathbf{v}$ can be calculated from the second equation of (3.12):

$$\Delta\mathbf{v} = \int \dot{\mathbf{v}} d\lambda = -2u^2 \int \nabla_{\perp}\Phi d\lambda. \quad (3.17)$$

One can then express deflection angle (3.16) as an integral over the background path as

$$\boldsymbol{\alpha} = 2 \int \nabla_{\perp}\Phi d\ell, \quad (3.18)$$

where $\ell = u\lambda$ is the spatial distance travelled along the unperturbed ray in the background spacetime.

To illustrate the application of (3.18), let us consider the gravitational deflection of light due to a point mass. We set up the Cartesian coordinates so that the background

path in Minkowski spacetime is along the x -direction. The impact parameter b is measured along the y -direction, which is the transverse direction from the path to the mass at the point of closest approach. The problem is effectively two-dimensional: one can always rotate the Cartesian axes $\{x, y, z\}$ so that the point mass, the ray, the light source and the observer all lie in the (x, y) -plane; we may therefore neglect the z -direction.

The Newtonian gravitational potential for a point mass is

$$\Phi = -\frac{M}{r} = -\frac{M}{\sqrt{x^2 + y^2}}, \quad (3.19)$$

where M is the mass of the gravitational source. The transverse gradient of the potential, evaluated on the background ray, is

$$\nabla_{\perp} \Phi = \frac{Mb}{(x^2 + b^2)^{3/2}} \mathbf{e}_y, \quad (3.20)$$

where \mathbf{e}_y is a unit vector in the y -direction. Assuming that the observer and source are located at $x \rightarrow \pm\infty$, the deflection angle (3.18) is

$$\boldsymbol{\alpha} = 2Mb \int_{-\infty}^{\infty} \frac{dx}{(x^2 + b^2)^{3/2}} \mathbf{e}_y = \frac{4M}{b} \mathbf{e}_y. \quad (3.21)$$

Transforming (3.21) to units in which the speed of light and Newton's gravitational constant are equal to c and G , respectively, we recover the expression (3.2) for the Einstein deflection angle.

An alternative approach to perturbative lensing involves considering the propagation of photons in a medium with effective refractive index $n = 1 - 2\Phi = 1 + 2|\Phi|$. The equations of motion for light rays can then be derived using Fermat's principle of least time [162, 165]. As in the case of standard ray optics, light travelling through a medium with a refractive index $n > 1$ travels slower than in vacuum.

3.3 Non-perturbative lensing

The observation of a range of gravitational lensing effects over the last forty years has played a crucial role in demonstrating the superiority of general relativity over the Newtonian theory. However, the mathematical formalism relies on quasi-Newtonian approximations, in which the gravitational field is assumed to be weak, and deflection angles are assumed to be small [70]; indeed, the largest of the observed lensing effects described in the previous section is only of the order of a few tens of arcseconds (see [166], for example).

Einstein's theory of general relativity leads to extreme gravitational objects, including neutron stars and black holes. Electromagnetic radiation propagates through the strong gravitational fields in the vicinity of such compact objects. It is therefore natural to want to build up an understanding of the signatures of *strong-field gravitational lensing*, associated with the near-field regions of extreme compact objects, where quasi-Newtonian approximations break down. In order to describe this strong-field gravitational lensing, we employ *non-perturbative lensing* [7], in which light rays follow null geodesics on a four-dimensional spacetime (\mathcal{M}, g_{ab}) , which is a solution to the field equations of general relativity.

3.3.1 Electromagnetism in curved spacetime

In its most general sense, gravitational lensing is concerned with all effects of a gravitational field (i.e., spacetime curvature) on the propagation of electromagnetic radiation. Here, we briefly review some fundamental aspects of Maxwell's theory of electromagnetism on curved spacetime.

Electromagnetism is governed by the *Faraday tensor* F_{ab} , which is antisymmetric type-(0, 2) tensor field (i.e., a two-form). At a particular spacetime point, the electric and magnetic fields are dependent on the choice of Lorentz frame. An observer moving on a timelike worldline with unit tangent vector u^a and orthonormal spatial frame $\{e_{(i)}^a\}$ would measure electric and magnetic fields with components $E_{(i)} = F_{ab} e_{(i)}^a u^b$ and $B_{(i)} = {}^*F_{ab} e_{(i)}^a u^b$, respectively. Here,

$${}^*F_{ab} = \frac{1}{2} \epsilon_{abcd} F^{cd} \quad (3.22)$$

is the Hodge dual of the Faraday two-form, where ϵ_{abcd} is the Levi-Civita tensor (see Section 2.1.3).

When describing electromagnetism in curved spacetime, it is convenient to introduce a complexified version of the Faraday tensor,

$$\mathcal{F}_{ab} = F_{ab} + i {}^*F_{ab}. \quad (3.23)$$

Since ${}^*\mathcal{F}_{ab} = -i\mathcal{F}_{ab}$, the complex bivector \mathcal{F}_{ab} is self-dual. We also note that $\mathcal{F}_{ab} \overline{\mathcal{F}}^{ab} = 0$, where an overbar denotes complex conjugation. From the complex bivector (3.23), one may define a complex three-vector \mathcal{F} whose components are given by $\mathcal{F}_{(i)} = \mathcal{F}_{ab} e_{(i)}^a u^b$; the real (imaginary) part of $\mathcal{F} = \mathbf{E} + i\mathbf{B}$ gives the electric (magnetic) field measured by an observer moving along the timelike worldline with tangent vector u^a with orthonormal spatial frame $\{e_{(i)}^a\}$.

In charge-free regions of spacetime, the Faraday two-form F_{ab} satisfies Maxwell's equations,

$$\nabla_b F^{ab} = 0, \quad \nabla_b {}^*F^{ab} = 0. \quad (3.24)$$

The second of these equations is equivalent to the Bianchi identity for the Faraday tensor: $\nabla_{[a} F_{cd]} = 0$. Hence, F is closed ($dF = 0$), so F is (locally) exact ($F = dA$, for some one-form potential A) by the Poincaré lemma; see Section 2.1.2. The Faraday tensor may be expressed in terms of the potential A_a as

$$F_{ab} = 2\nabla_{[a} A_{b]}. \quad (3.25)$$

Expanding out the covariant derivatives in (3.25), we see that the connection coefficients will cancel by antisymmetry; this permits us to write $F_{ab} = 2\partial_{[a} A_{b]}$. The fact that the Faraday tensor is locally exact means that it is invariant under gauge transformations of the form $A \mapsto A' = A + dG$, where G is any scalar field.

In the absence of charges, one may express Maxwell's equations in terms of the complexified Faraday tensor (3.23) as

$$\nabla_b \mathcal{F}^{ab} = 0, \quad \nabla_{[a} \mathcal{F}_{bc]} = 0. \quad (3.26)$$

It is possible to derive a wave equation for the Faraday tensor, by taking a derivative of the first of Maxwell's equations (3.24), applying the commutator identity for covariant derivatives, and the Bianchi identity. This yields

$$\square F_{ab} + 2R_{abcd}F^{cd} + R_a{}^c F_{bc} - R_b{}^c F_{ac} = 0, \quad (3.27)$$

where $\square = \nabla^a \nabla_a$ is the d'Alembertian operator. In (3.27), one may replace the Faraday tensor F_{ab} with its complexified counterpart \mathcal{F}_{ab} , as we only consider the source-free Maxwell equations (3.24).

One may instead derive a wave equation for the vector potential: in the absence of charges, this reads

$$\square A^a - R^a{}_b A^b - \nabla^a (\nabla_b A^b) = 0. \quad (3.28)$$

Adopting the Lorenz gauge condition $\nabla_a A^a = 0$ eliminates the final term on the left-hand side of (3.28).

When expressed in terms of the complexified Faraday tensor (3.23), the electromagnetic stress–energy tensor (see Section 2.3.2) takes a particularly simple form:

$$T_{ab} = \frac{1}{4\pi} \left(F_{ac} F_b{}^c - \frac{1}{4} g_{ab} F_{cd} F^{cd} \right) = \frac{1}{8\pi} \text{Re}(\mathcal{F}_a{}^c \overline{\mathcal{F}}_{bc}). \quad (3.29)$$

Clearly, the stress–energy tensor is trace-free: $T^a{}_a = \frac{1}{2} \mathcal{F}_{ab} \overline{\mathcal{F}}^{ab} = 0$. Furthermore, in the absence of electromagnetic sources, the stress–energy tensor (3.29) satisfies the conservation equation $\nabla_b T^{ab} = 0$.

3.3.2 Geometric optics

Leading-order geometric optics approximation

Consider an electromagnetic field propagating in curved spacetime. In particular, let us assume that the wavelength is short in comparison with other physically relevant length scales, and that the inverse frequency is short in comparison with relevant time scales. In such situations, one may introduce a geometric optics ansatz for the vector potential A_a , and insert this into the wave equation (3.28), restricting to Lorenz gauge, $A^a{}_{;a} = 0$ [29, 167].

Alternatively, one may introduce a geometric optics ansatz for the Faraday tensor itself (or its complexified counterpart) [168–170]. This is particularly useful when calculating the stress–energy tensor (3.29); moreover, results obtained by working with the Faraday tensor are manifestly gauge-invariant, in contrast to those obtained from the vector potential A^a . Here, we summarise the key points of the formalism presented by Dolan [81, 171] for the electromagnetic field; similar approaches can be taken for scalar and gravitational fields [171].

In this approach, we introduce the geometric optics ansatz for (3.23) of the form

$$\mathcal{F}_{ab} = \mathcal{A} \mathfrak{f}_{ab} \exp(i\omega\Phi), \quad (3.30)$$

where ω is an order-counting parameter which is related to the frequency; \mathcal{A} is the (real) amplitude; Φ is the (real) phase; and \mathfrak{f}_{ab} is the polarisation bivector, a complex self-dual bivector field. Substitution of the ansatz (3.30) into the wave equation for \mathcal{F}_{ab} gives

$$-\omega^2 k^c k_c \mathfrak{f}_{ab} + i\omega \left[(2k^c \mathcal{A}_{;c} + k^c{}_{;c} \mathcal{A}) \mathfrak{f}_{ab} + \mathcal{A} k^c \mathfrak{f}_{ab;c} \right] + O(\omega^0) = 0, \quad (3.31)$$

where $k^a = g^{ab}\nabla_b\Phi$ is the gradient of the phase.

At leading order, we see from (3.31) that k^a is a null vector field ($k^a k_a = 0$). Since this vector field is both null and a gradient, it is quick to show that k^a satisfies the geodesic equation, $k^b k_{a;b} = 0$. In other words, the integral curves of k^a – spacetime paths¹ $x^a(s)$ with $k^a = \dot{x}^a$ – are null geodesics which lie in constant-phase hypersurfaces ($\Phi = \text{constant}$); these are known as *null generators*, and may be obtained from the Hamiltonian $H = \frac{1}{2}g^{ab}k_a k_b$, where the “conjugate momenta” are $k_a = g_{ab}\dot{x}^b$, and the Hamiltonian function H vanishes along null geodesics. Furthermore, it follows directly from the fact that $k_a = \nabla_a\Phi$, and the equality of mixed partial derivatives, that $\nabla_{[a}k_{b]} = \nabla_{[a}\nabla_{b]}\Phi = 0$. Hence, the integral curves of k^a form a *twist-free* null congruence.

Proceeding to $O(\omega)$, we see from (3.31) that we require $(2k^c\mathcal{A}_{;c} + k^c{}_{;c}\mathcal{A})\mathfrak{f}_{ab} + \mathcal{A}k^c\mathfrak{f}_{ab;c} = 0$. Exploiting the ambiguity in the definition of the amplitude and polarisation bivector in (3.30), we split the equations such that

$$k^a\mathcal{A}_{;a} = -\frac{1}{2}\vartheta\mathcal{A}, \quad (3.32)$$

$$k^c\mathfrak{f}_{ab;c} = 0, \quad (3.33)$$

where $\vartheta = k^a{}_{;a}$ is the *expansion scalar* [172]. By (3.32), the flux $\mathcal{A}^2 k^a$ is conserved: $\nabla_a(\mathcal{A}^2 k^a) = 0$. The transport equation (3.33) says that the polarisation bivector \mathfrak{f}_{ab} is parallel-transported along the null generators. Moreover, at leading order, Maxwell’s equation $\nabla_b\mathcal{F}^{ab} = 0$ implies that the polarisation bivector is transverse: $\mathfrak{f}_{ab}k^b = 0$.

A circularly polarised wave which satisfies $\mathfrak{f}_{ab}k^b = 0$ and $k^c\mathfrak{f}_{ab;c} = 0$ can be constructed by choosing

$$\mathfrak{f}_{ab} = 2k_{[a}m_{b]}, \quad (3.34)$$

where k^a is the gradient of the phase, and m^a is a complex null ($m^a m_a = 0$) vector field which is transverse ($k^a m_a = 0$) and satisfies $m^a \bar{m}_a = 1$. In addition, m^a is required to satisfy the transport equation $k^b m^a{}_{;b} = \alpha(s)k^a$, where $\alpha(s)$ is a scalar function along the null generator. We have $\alpha = 0$ in the case where m^a is parallel-transported along the null generators. (We note that a linearly polarised wave may be constructed by superposing two circularly polarised waves [171].)

The complex null vector m^a may be constructed from the legs of an orthonormal triad $\{e_{(i)}^a\}$, viz. $m^a = \frac{1}{\sqrt{2}}(e_{(1)}^a + ie_{(2)}^a)$; see Section 2.5.1. The handedness of the circularly polarised wave depends on the sign of ω and the handedness of m^a . We hereafter assume that m^a is constructed such that $i\epsilon_{abcd}u^a k^b m^c \bar{m}^d > 0$ for any future-pointing timelike vector u^a . The wave is then right-hand (left-hand) polarised if $\omega > 0$ ($\omega < 0$). We note that the Lorentz transformation $m^a \mapsto e^{i\Theta}m^a + Bk^a$, where Θ is real and B is complex, preserves the handedness, transversality property, and the parallel-transport of m^a (see Section 2.5.2).

One may introduce an auxiliary null vector n^a , which is future-pointing, and satisfies $k^a n_a = -1$ with all other inner products zero. The tetrad $\{k^a, n^a, m^a, \bar{m}^a\}$ is then a complex null tetrad (see Section 2.5.1): the Newman–Penrose formalism is therefore well-adapted to describe the geometric optics approximation for the electromagnetic field on

¹Here, an overdot denotes differentiation with respect to the affine parameter s .

curved spacetime. It is convenient to choose this tetrad to be parallel-propagated along null geodesics with tangent vector field k^a .

Inserting (3.30) and (3.34) into (3.29), we find that the electromagnetic stress–energy tensor takes the form of a null fluid at leading order:

$$T_{ab} = \frac{1}{8\pi} \mathcal{A}^2 k_a k_b. \quad (3.35)$$

Conservation of the stress–energy tensor implies that the flux is conserved at leading order in ω , i.e., $\nabla_a (\mathcal{A}^2 k^a) = 0$.

Gravitational lensing from geometric optics

In order to calculate many quantities of interest to quantify gravitational lensing phenomena (e.g. image distortion and brightness), one must solve the geodesic deviation equation (or Jacobi equation), which describes how spacetime curvature leads to a relative acceleration between neighbouring geodesics in a congruence. The geodesic deviation equation was derived for a general geodesic congruence in Section 2.2.4. We recall here that, for a *null* geodesic congruence $x^a(s, v)$ with tangent vector field $k^a = \frac{\partial x^a}{\partial s}$ and separation vector $\xi^a = \frac{\partial x^a}{\partial v}$, the geodesic deviation equation is

$$D^2 \xi^a = -R^a{}_{bcd} k^b \xi^c k^d, \quad (3.36)$$

where $D = k^a \nabla_a$ denotes the directional derivative along the tangent vector field. Recall from Section 2.2.4 that $\mathcal{L}_k \xi = 0$ implies that the quantity $\xi^a k_a$ is conserved along each null ray. If we choose ξ^a to be initially transverse to k^a ($\xi^a k_a = 0$), then it will remain so along the ray.

The null geodesic deviation equation (3.36) describes the evolution of an infinitesimal bundle of rays with elliptical cross-section along the central null geodesic $x^a(s, 0)$. In this context, an *infinitesimal bundle of rays* is the set [7]

$$\mathcal{B} = \{ \alpha_1 \xi_1^a + \alpha_2 \xi_2^a \mid \alpha_1, \alpha_2 \in \mathbb{R}, \alpha_1^2 + \alpha_2^2 \leq 1 \}, \quad (3.37)$$

where ξ_1 and ξ_2 are two vector fields which satisfy (3.36) and $g_{ab} k^a \xi^b = 0$, such that $\xi_1(s)$, $\xi_2(s)$ and $k^a(s)$ are linearly independent for almost all values of the affine parameter s [7].

A standard approach [7] to solving (3.36) is to introduce a Sachs basis, then decompose (3.36) into the Sachs equations for the optical scalars [160, 173]. Here, we favour a slightly different approach, which involves working with the (parallel-propagated) complex null tetrad $\{k^a, n^a, m^a, \bar{m}^a\}$ introduced above. The deviation vector ξ^a , restricted to the central null geodesic, may be decomposed in terms of the complex null tetrad along that geodesic:

$$\xi^a = a(s)k^a + b(s)n^a + \bar{c}(s)m^a + c(s)\bar{m}^a, \quad (3.38)$$

where a and b are real functions along the ray, and c is complex. Substitution of (3.38) into (3.36) yields a hierarchical system of coupled second-order ordinary differential equations [81, 171], namely

$$\ddot{b} = 0, \quad (3.39)$$

$$\ddot{a} = bR_{abcd}k^a n^b k^c n^d + cR_{abcd}k^a n^b k^c \bar{m}^d + \bar{c}R_{abcd}k^a n^b k^c m^d, \quad (3.40)$$

$$\ddot{c} = -bR_{abcd}k^a m^b k^c n^d - cR_{abcd}k^a m^b k^c \bar{m}^d - \bar{c}R_{abcd}k^a m^b k^c m^d. \quad (3.41)$$

The first of these equations is consistent with the fact that $\xi^a k_a = -b$ is conserved. Setting $b = 0$ ensures that ξ^a is transverse to k^a . The equations then read

$$\ddot{a} = (\bar{\Phi}_{00} + \bar{\Psi}_1) c + (\Phi_{00} + \Psi_1) \bar{c}, \quad (3.42)$$

$$\ddot{c} = -\Phi_{00} c - \Psi_0 \bar{c}, \quad (3.43)$$

where $\Phi_{00} = \frac{1}{2} R_{ab} k^a k^b$ is a complex Ricci scalar, and $\Psi_0 = C_{abcd} k^a m^b k^b m^d$ and $\Psi_1 = C_{abcd} k^a n^b k^b m^d$ are complex Weyl scalars; see Section 2.5 for a review of the Newman–Penrose formalism.

Applying the directional derivative D to $\xi^a m_a$, and using the definition (3.38), as well as the Lie-transport ($\mathcal{L}_k \xi^a = 0$) and transversality ($k^a \xi_a = -b = 0$) properties of the deviation vector, we find $\dot{c} = -\rho c - \sigma \bar{c}$, where $\rho = -m^a k_{a;b} \bar{m}^b$, and $\sigma = -m^a k_{a;b} m^b$ are Newman–Penrose scalars. Inserting this into (3.43) and equating coefficients of c and \bar{c} yields the first-order transport equations

$$\dot{\rho} = \rho^2 + \sigma \bar{\sigma} + \Phi_{00}, \quad (3.44)$$

$$\dot{\sigma} = \sigma(\rho + \bar{\rho}) + \Psi_0, \quad (3.45)$$

which are known as the *Sachs equations* [160, 173]. The real and imaginary parts of $-\rho = \theta + i\varpi$ and $-\sigma = \varsigma_1 + i\varsigma_2$ are known collectively as the *optical scalars* [173–175]; in particular, $\theta = \frac{1}{2}\vartheta = \frac{1}{2}k^a{}_{;a}$ is the *expansion*, ϖ is the *twist*, and $(\varsigma_1, \varsigma_2)$ is the *shear*. In the case of a hypersurface-orthogonal null congruence (e.g. the geometric optics congruence where $k_a = \nabla_a \Phi$ is orthogonal to constant-phase hypersurfaces $\Phi = \text{constant}$), we have $\varpi = 0$. Such a congruence is called *twist-free*.

One limitation of this approach is that the Sachs equations (3.44)–(3.45) suffer from divergences at *caustic points*, where neighbouring rays in the congruence cross [7, 81, 171]; we revisit this issue in Chapter 7 in the context of (higher-order) geometric optics on Kerr spacetime. The second-order equation (3.43) does not exhibit the same pathological behaviour at caustic points. One may determine the optical scalars (ρ, σ) by solving (3.43), then inverting the matrix equation

$$\begin{bmatrix} \dot{c}_1 \\ \dot{c}_2 \end{bmatrix} = - \begin{bmatrix} c_1 & \bar{c}_1 \\ c_2 & \bar{c}_2 \end{bmatrix} \begin{bmatrix} \rho \\ \sigma \end{bmatrix}, \quad (3.46)$$

where c_1 and c_2 are any linearly independent pair of solutions to the second-order equation (3.43). We note that the matrix on the right-hand side of (3.46) is non-invertible at points where $\text{Im}(c_1 \bar{c}_2) = 0$, which are precisely the caustic points.

Using a similar argument to that presented above, one may show that $\dot{a} = \bar{\tau} c + \tau \bar{c}$, where $\tau = -m^a k_{a;b} n^b$ is a Newman–Penrose scalar. One may find τ by inverting

$$\begin{bmatrix} \dot{a}_1 \\ \dot{a}_2 \end{bmatrix} = \begin{bmatrix} c_1 & \bar{c}_1 \\ c_2 & \bar{c}_2 \end{bmatrix} \begin{bmatrix} \bar{\tau} \\ \tau \end{bmatrix}, \quad (3.47)$$

for any linearly independent pair of solutions (a_1, c_1) and (a_2, c_2) to the second-order system of ordinary differential equations (3.42)–(3.43). Again, this breaks down at caustic points, where neighbouring rays in the bundle cross.

Let $e_{(1)}^a$ and $e_{(2)}^a$ be orthogonal unit vectors which lie in the instantaneous electromagnetic wavefront, i.e., they are orthogonal to k^a . The complex number $c = \frac{1}{\sqrt{2}}(x + iy)$ corresponds to a point in the instantaneous wavefront with position vector $\hat{\xi}^a = \bar{c}m^a + c\bar{m}^a = xe_{(1)}^a + ye_{(2)}^a$, where we recall the definition of the complex null vector $m^a = \frac{1}{\sqrt{2}}(e_{(1)}^a + ie_{(2)}^a)$; and we note that $e_{(1)}^a = \frac{1}{\sqrt{2}}(m^a + \bar{m}^a)$ and $e_{(2)}^a = -\frac{i}{\sqrt{2}}(m^a - \bar{m}^a)$. If c_1 and c_2 are any pair of linearly independent solutions to (3.43), then the function $c(\phi) = \cos \phi c_1 + \sin \phi c_2$ parametrises an infinitesimal bundle of rays with elliptical cross-section which lies in the wavefront. One can demonstrate that the principal axes of the ellipse are given by $c_+ = \cos \phi_0 c_1 + \sin \phi_0 c_2$ and $c_- = -\sin \phi_0 c_1 + \cos \phi_0 c_2$, with $\tan(2\phi_0) = \frac{2\text{Re}(c_1\bar{c}_2)}{|c_1|^2 - |c_2|^2}$. The semi-major and semi-minor axes are $d_{\pm} = \sqrt{2}|c_{\pm}|$, with $d_+d_- = 2|\text{Im}(c_1\bar{c}_2)|$ and $d_+^2 + d_-^2 = 2(|c_1|^2 + |c_2|^2)$. The parameters d_+ and d_- are sometimes referred to as the *shape parameters* of the bundle [7]. The cross-sectional area of the infinitesimal bundle of rays is $A = \pi d_+d_-$; this quantity satisfies the transport equation

$$DA = -(\rho + \bar{\rho})A = \vartheta A, \quad (3.48)$$

where $\rho = -m^a k_{a;b} \bar{m}^b$ is a Newman–Penrose scalar and ϑ is the expansion scalar, cf. (3.32).

Using a derivation analogous to that of the geodesic deviation equation (Section 2.2.4), Dolan [81, 171] also derives the *differential precession equation*,

$$D\zeta^a = -R^a{}_{bcd}m^b\bar{\xi}^ck^d, \quad (3.49)$$

where $\zeta^a = \xi^b \nabla_b m^a|_C$ is the *precession vector*, defined on an observer's worldline C . Projecting on the null tetrad, $\zeta^a = v(s)k^a + \dot{c}(s)n^a + w(s)m^a$, yields a system of first-order ordinary differential equations for the coefficients v and w which read

$$\dot{v} = cR_{abcd}n^am^b\bar{m}^ck^d + \bar{c}R_{abcd}n^am^bm^ck^d, \quad (3.50)$$

$$\dot{w} = -cR_{abcd}\bar{m}^am^b\bar{m}^ck^d - \bar{c}R_{abcd}\bar{m}^am^bm^ck^d. \quad (3.51)$$

In a Ricci-flat spacetime, this system of equations is simply

$$\dot{v} = \bar{\Psi}_2 c, \quad (3.52)$$

$$\dot{w} = -\bar{\Psi}_1 c - \Psi_1 \bar{c}, \quad (3.53)$$

where $\Psi_2 = C_{abcd}k^am^b\bar{m}^cn^d$ is a complex Weyl scalar. Taking a derivative of ζ^a contracted with the legs of the complex null tetrad, one may show that $v = \bar{\mu}c + \bar{\lambda}\bar{c}$ and $w = -\bar{\chi}c + \chi\bar{c}$, where μ , λ and $\chi = \beta - \bar{\alpha}$ are Newman–Penrose scalars (Section 2.5.3). The Newman–Penrose quantities μ , λ and χ can therefore be obtained by inverting

$$\begin{bmatrix} v_1 \\ v_2 \end{bmatrix} = \begin{bmatrix} \bar{c}_1 & c_1 \\ \bar{c}_2 & c_2 \end{bmatrix} \begin{bmatrix} \bar{\mu} \\ \bar{\lambda} \end{bmatrix}, \quad \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} = \begin{bmatrix} \bar{c}_1 & c_1 \\ \bar{c}_2 & c_2 \end{bmatrix} \begin{bmatrix} \chi \\ -\bar{\chi} \end{bmatrix}, \quad (3.54)$$

for any pair of linearly independent solutions (c_1, v_1, w_1) and (c_2, v_2, w_2) to the equations (3.43), (3.52) and (3.53). This inversion breaks down at caustic points; however, the first-order ordinary differential equations (3.52)–(3.53) can be evolved through caustic points without issue, as in the case of (3.42)–(3.43).

3.3.3 Photon orbits

In general relativity, ultra-compact objects – such as black holes – are responsible for spacetime curvature, which can cause extreme local light deflection. In fact, such objects (by definition) possess light-rings (i.e., circular photon orbits), and are able to bend light through an arbitrarily large angle [60]. We consider null geodesic motion on Kerr spacetime as an illustrative example. (This is considered in more detail in Chapters 6 and 7, so we only outline the key details here.)

The Kerr solution, given in Boyer–Lindquist coordinates $\{t, r, \theta, \phi\}$ in (2.86), describes the exterior spacetime of a rotating black hole of mass M with spin a (see Section 2.4 for a review). The spacetime admits a pair of commuting Killing vectors $\xi_{(t)} = \partial_t$ and $\xi_{(\phi)} = \partial_\phi$, which correspond to the stationarity and axisymmetry of the geometry. In Boyer–Lindquist coordinates, the geodesic motion is Liouville integrable thanks to the existence of the principal tensor h_{ab} , from which one can generate the Killing tensor K_{ab} (see Section 2.6.4). Along a null geodesic with tangent vector field k^a , the Hamiltonian $H = \frac{1}{2}g^{ab}k_a k_b$, total energy $E = -\xi_{(t)}^a k_a$, azimuthal angular momentum $L_z = \xi_{(\phi)}^a k_a$, and Carter constant $K = K^{ab}k_a k_b$ are all conserved, with $H = 0$ in the null case.

Null geodesics in the equatorial plane ($\theta = \frac{\pi}{2}$), for which $K = (aE - L_z)^2$, are governed by the radial equation [176]

$$\dot{r} = \frac{L_z^2}{b^2} \left[1 + \frac{(a^2 - b^2)}{r^2} + \frac{2M(b - a)^2}{r^3} \right], \quad (3.55)$$

where an overdot denotes differentiation with respect to an affine parameter, and b is an *impact parameter*, defined as the ratio of the azimuthal angular momentum to the energy:

$$b = \frac{L_z}{E}. \quad (3.56)$$

Unstable circular photon orbits, which neither plunge into the black hole nor escape to spatial infinity, are possible if the pair of conditions $\dot{r} = 0 = \ddot{r}$ are met. Prograde (+) and retrograde (−) perpetual orbits exist at Boyer–Lindquist radii [176]

$$\hat{r}_\pm = 2M \left[1 + \cos \left(\frac{2}{3} \arccos \left(\mp \frac{a}{M} \right) \right) \right]. \quad (3.57)$$

The corresponding critical impact parameters are given by

$$b_\pm = \pm 3\sqrt{M\hat{r}_\pm} - a, \quad (3.58)$$

where $b_- < 0$ by definition for the retrograde orbit. In the literature, the circular photon orbits of constant Boyer–Lindquist radii are referred to as *unstable light-rings*. In the non-rotating case ($a = 0$), the Kerr spacetime reduces to the Schwarzschild solution, in which there exists a unique unstable light-ring at $r = 3M$, with critical impact parameter $b = 3\sqrt{3}M$.

The null orbits around a Schwarzschild black hole are necessarily confined to a plane passing through the black hole’s centre (which can be taken to be the equatorial plane after a suitable change of coordinates), due to the spherical symmetry of the spacetime

geometry. In Kerr spacetime, however, the lack of spherical symmetry means that non-planar orbits are permitted. In fact, there exists a set of so-called *spherical photon orbits* – unstable perpetual null geodesics of constant Boyer–Lindquist radii – which are not necessarily confined to the plane $\theta = \frac{\pi}{2}$ [76, 177]. The subfamily of spherical photon orbits which are restricted to the equatorial plane are precisely the two light-rings located at $r = \hat{r}_{\pm}$.

In stationary axisymmetric spacetimes whose motion is not (necessarily) Liouville integrable, there may exist generalisations of the spherical photon orbits which are present in Kerr spacetime. In recent work, Cunha *et al.* [79] have generalised the notion of a spherical photon orbit to a generic stationary axisymmetric spacetime (\mathcal{M}, g_{ab}) . Let $\gamma(s)$ be an affinely parametrised null geodesic. Then $\gamma(s)$ is a *fundamental photon orbit* if it is restricted to a compact spatial region (i.e., it is a bound state), and there exists a value $T > 0$ such that $\gamma(s) = \gamma(s + T)$ for all $s \in \mathbb{R}$ (i.e., it is periodic with period T).

Fundamental photon orbits on Kerr spacetime are precisely the spherical photon orbits, which include the equatorial light-rings as a subset. In the exterior region (i.e., outside the outer event horizon), all spherical photon orbits are unstable. More generally, however, fundamental photon orbits can be *stable*, leading to interesting phenomenological features, such as spacetime instabilities. This will be central to the work presented in Chapter 6.

3.3.4 Black hole shadows

Black holes are characterised by their event horizon, a one-way causal boundary in spacetime beyond which nothing – not even light – can escape once it has fallen inside (see Section 2.4.2). For an observer in spacetime, the *black hole shadow* is the apparent optical image of the black hole. In a given observation frame, the black hole shadow corresponds to the set of all null directions in the local sky which originate (asymptotically) from the black hole’s event horizon: the shadow thus corresponds to a lack of radiation received by the observer. In the high-frequency (geometric optics) approximation, the shadow is associated with the black hole’s light absorption cross-section. Moreover, the outline of the shadow is intimately related to the gravitational lensing of electromagnetic radiation in the strong-field region; the shadow therefore encodes important information about the geometry of spacetime close to the black hole.

To illustrate the concept of a black hole shadow, consider the case of null geodesics on Schwarzschild spacetime. The unstable light-ring at $r = 3M$ plays an important role: this orbit defines the boundary between the null rays which plunge into the event horizon of the black hole, and those which escape to infinity. The shadow of the Schwarzschild black hole is simply a circular disc in the sky, whose radius is determined by the unstable light-ring. For a static observer in Schwarzschild spacetime, the angular radius of the black hole shadow was calculated by Synge [72]. (In fact, Synge calculated the “escape cone” of light, which is simply the complement of the shadow on the observer’s local sky.)

In 1979, Luminet [73] determined the image of a “bare” Schwarzschild black hole, illuminated by a distant light source from directly behind. In the same paper, Luminet also analysed the more astrophysically relevant case of a Schwarzschild black hole surrounded by an emitting accretion disc. In situations where one is only concerned with the outline of

the shadow, the system (i.e., the black hole and the observer) is assumed to be surrounded by a *celestial sphere* at infinity, which acts as a light source.

Kerr black hole shadow

Here we consider the shadow of a Kerr black hole of mass M and spin a (see Section 2.4). For a null geodesic with energy E , angular momentum L_z , and Carter constant K , the motion is determined by the independent *impact parameters* [74]

$$b = \frac{L_z}{E}, \quad \chi = \frac{K - (aE - L_z)^2}{E^2}. \quad (3.59)$$

The shadow of the Kerr black hole is determined by the spherical photon orbits of constant Boyer–Lindquist radius r (satisfying $\dot{r} = 0 = \ddot{r}$). These spherical photon orbits separate the rays which are captured by the black hole from those which escape to infinity. The former correspond to the black hole shadow; the latter correspond to regions of the observer’s local sky which are illuminated by light from the celestial sphere. For a spherical photon orbit with Boyer–Lindquist radius r , the impact parameters (3.59) are [76]

$$b = \frac{r^3 - 3Mr^2 + a^2r + Ma^2}{a(M - r)}, \quad (3.60)$$

$$\chi = \frac{r^2}{r^2 - a^2} (3r^2 + a^2 - b^2). \quad (3.61)$$

Consider a static observer at infinity with a *viewing angle* of $\theta = \theta_0$ in Boyer–Lindquist coordinates. The boundary of the shadow is the locus of points in the observer’s local sky which correspond to null rays that asymptote to the spherical photon orbits; the orbital parameters are therefore determined by (3.60) and (3.61). The observer will set up an *image plane* which has local Cartesian-type coordinates (X, Y) , which are related to the impact parameters (b, χ) by [55, 60, 74]

$$X = -b \operatorname{cosec} \theta_0, \quad Y = \pm \sqrt{\chi + a^2 \cos^2 \theta_0 - b^2 \cot^2 \theta_0}. \quad (3.62)$$

The shadow boundary is then defined as a parametric curve in the (X, Y) -plane, where X and Y depend on the Boyer–Lindquist radius r of a spherical photon orbit through the parameters b and χ . The shadow boundary demarcates all of the rays which reach the observer from those which plunge into the black hole. Recently, Cunha and Herdeiro [60] provided an expression for the shadow boundary as a function $Y(X)$, rather than in the parametric form (3.62).

In Figure 3.1, we show the shadow boundary of a Kerr black hole of mass M and spin parameter a , for an observer located in the equatorial plane ($\theta = \frac{\pi}{2}$). Figure 3.1(a) shows the Schwarzschild case ($a = 0$); the shadow boundary is a circle whose radius in the image plane is determined by the equatorial unstable photon orbit (light-ring) at $r = 3M$ with impact parameter $b = 3\sqrt{3}M$. As the value of a is increased from zero, as shown in Figures 3.1(b) and 3.1(c), the frame-dragging effects are clearly visible: the shadow is asymmetric, due to the fact that $|b_-| > |b_+|$ (i.e., there is more absorption of retrograde rays than prograde rays); and it is deformed from circularity, exhibiting a characteristic “D”-shape in the extremal ($a = M$) case.




Figure 3.1. Boundary of a Kerr black hole shadow for a selection of values of the spin parameter a , as seen by a distant observer with a viewing angle of $\theta = \frac{\pi}{2}$ and image plane coordinates (X, Y) . In the Schwarzschild case ($a = 0$), the shadow boundary is a circle of radius $3\sqrt{3}M$; the shadow is entirely determined by the unstable photon orbit at $r = 3M$ with impact parameter $b = 3\sqrt{3}M$, due to the spherical symmetry of the spacetime. As one increases the value of a , the shadow boundary becomes deformed and asymmetric, due to the frame-dragging effects caused by the rotation of the black hole.

Ray-tracing

Often it is impractical (or impossible) to calculate the shadow boundary analytically in terms of non-escaping photon orbits, e.g. in situations where the geodesic motion is non-integrable. In such cases, the null geodesic equations must be solved using numerical methods. Rather than evolve all light rays which begin on the celestial sphere and determine which of these rays reach the observer, it is computationally more effective to reverse this process, tracing rays away from the observer, backwards in time. In this approach, known as *backwards ray-tracing*, the black hole shadow corresponds to the set of all null geodesics which asymptote to the black hole's event horizon when traced backwards in time from the observer.

The use of ray-tracing in general relativity dates back to 1970s, when Cunningham and Bardeen [178] determined the optical appearance of a star orbiting an extremal Kerr black hole, Cunningham [179] calculated the emission spectrum of an accretion disc around a Kerr black hole, and Luminet [73] determined the image of an accretion disc around a Schwarzschild black hole.

Since then, a variety of numerical ray-tracing codes have been developed, including GeoKerr [180], GeodesicViewer [181], GYOTO [182], GeoViS [183], and PyHole [184]. Ray-tracing has been employed in a range of situations to understand gravitational lensing phenomena, such as the optical appearance of the black hole (i.e., the shadow) [9, 80]; images of accretion structures around black holes [180, 185]; the apparent appearance of stars orbiting compact objects [186, 187]; magnetohydrodynamical processes occurring around black holes [188]; the polarisation of radiation propagating in the vicinity of compact objects [189]; the shadow of a dynamical binary black hole system [190]; and the lensing dynamics of hairy black holes [191].

Chapter 4

Binary black hole shadows and chaotic scattering

4.1 Introduction

Pairs of rotating black holes that orbit one another – known as binary black holes – have recently been detected. In 2015, the first direct detection of gravitational waves was made by the Advanced LIGO experiment, which confirmed the existence of binary black holes in nature [14]. Since then, a host of gravitational-wave signals from merging binary black holes have been observed by the LIGO–Virgo collaboration [17, 192–195]. More recently, the era of multi-messenger astronomy began with the observation of a gravitational-wave signal from binary neutron star merger [15], which was accompanied, approximately 1.7 seconds later, by a gamma-ray burst [16]. A catalogue of compact binary coalescences observed by LIGO and Virgo during the first and second observing runs can be found in [17, 18].

The Event Horizon Telescope (EHT) – a global array of radio telescopes, which employs very-long-baseline interferometry (VLBI) to create an Earth-scale virtual telescope – began to observe nearby galactic centres, including the Milky Way and Messier 87 (M87), in 2017 [23, 80, 196, 197]. The EHT’s principal aim is to image the environment of supermassive black hole candidates which reside at the centre of these galaxies, with the hope of observing gravitational lensing phenomena associated with the black hole, such as multiple images, Einstein rings, and the black hole shadow itself. On 10 April 2019, the EHT collaboration presented the first event-horizon-scale image of the supermassive black hole candidate M87* (at the centre of the galaxy M87) from VLBI observations conducted in April 2017 [1–6]. The image captured by the EHT is consistent with expectations for the shadow of a Kerr black hole, as predicted by general relativity (see Section 3.3.4).

Inspired by the recent detection of gravitational waves from merging binary black holes, and the first direct image of a black hole’s event horizon, a strand of recent work has focussed on the strong-field gravitational lensing effects of binary black holes. The geometry of a dynamical binary, such as those observed by the LIGO–Virgo collaboration, is not known analytically as a solution to Einstein’s field equations. In order to analyse the lensing effects of such systems, one must perform ray-tracing on top of fully non-linear

numerical simulations. Bohn *et al.* [190] consider what an observer in the vicinity of a binary black hole system would actually *see* as the two black holes orbit, inspiral and merge. The authors find that a binary black hole shadow is not simply the superposition of two singleton black hole shadows. In addition, the binary shadow exhibits intricate structures, including “eyebrow” features – arc-shaped shadows which surround the primary globular shadow. Such features were anticipated by Nitta *et al.* [198] and Yumoto *et al.* [199] in their studies of colliding black holes using the cosmological Kastor–Traschen solution [200].

Analysing the lensing effects of dynamical binary black holes is a formidable task, because it requires computationally expensive numerical simulations of the field equations of general relativity. In this chapter, based on the work presented in [201], our principal aim is to understand the *qualitative* features of binary black hole shadows. We therefore choose to focus on an imitative closed-form model: the Majumdar–Papapetrou binary black hole solution [202, 203], a static solution to the Einstein–Maxwell equations which describes a pair of extremally charged (Reissner–Nordström) black holes whose gravitational attraction and electrostatic repulsion are in perfect balance. The Majumdar–Papapetrou solution has been studied in detail by Hartle and Hawking [204] and Chandrasekhar [89].

In Section 2.6.4, we saw that a singleton Kerr–Newman black hole admits a Killing tensor, which gives rise to a fourth “hidden” constant of geodesic motion (Carter’s constant), permitting the separability and Liouville integrability of the geodesic equations. By contrast, the presence of a pair of black holes in the Majumdar–Papapetrou binary black hole system reduces the symmetry, eliminating the Killing tensor associated with the Carter constant. As a result, the null geodesic equations, which govern the motion of photons on the Majumdar–Papapetrou geometry, are non-integrable; we should therefore anticipate rich gravitational lensing phenomena around binary black holes.

Here, building on the work of Contopoulos [140–142, 205–208] and others [127, 134, 138, 139, 209–218], we advance the view that binary black hole systems are natural examples of *chaotic scattering*: the system admits (we shall show) a fractal set of scattering singularities of measure zero. (See Section 2.6.5 for a review of chaotic scattering in Hamiltonian systems.) In regions of chaotic scattering, a small uncertainty in fixing the initial conditions can lead to a drastically different outcome in the final state of the system. For the Majumdar–Papapetrou binary black hole system, the scattering singularities are associated with *perpetual orbits*, i.e., unstable null geodesics which are neither scattered nor absorbed by the black holes.

This chapter is based on [201]. Section 4.2 contains a review of the two-centre Majumdar–Papapetrou solution. In Section 4.3, we discuss the null geodesic structure of the Majumdar–Papapetrou spacetime, with a focus on the Lagrangian and Hamiltonian formalisms, and a review of integrability and chaotic motion of null geodesics. Next, in Section 4.4, we analyse a planar scattering problem and its associated one-dimensional binary black hole shadows. In Section 4.5, we introduce symbolic dynamics, considering the Gaspard–Rice three-disc scatterer as an illustrative example; and use the technique to describe the null geodesics of the Majumdar–Papapetrou binary black hole. In Section 4.6, we use symbolic dynamics to understand the ordering of the scattering singularities in the initial data; we

construct the one-dimensional binary black hole shadows using an iterative procedure akin to that of the Cantor set, based on symbolic dynamics; we demonstrate self-similarity of the one-dimensional shadow explicitly; and we discuss quantitative measures of chaos which can be employed to understand the fractal structure of the strange repeller (see Section 2.6.5). In Section 4.7, we analyse non-planar null geodesics, which can be understood using an effective null geodesic potential; we highlight the existence of stable bounded photon orbits; and we present ray-traced images of two-dimensional binary black hole shadows for the Majumdar–Papapetrou system. In Section 4.8, we consider the consequences of following rays through the event horizons in the maximally extended Majumdar–Papapetrou binary black hole spacetime. Finally, we conclude with a discussion of the work presented in this chapter in Section 4.9.

4.2 Majumdar–Papapetrou solution

The Majumdar–Papapetrou spacetime is a solution to the Einstein–Maxwell equations of gravity and electromagnetism in electrovacuum. The geometry is described in isotropic coordinates $\{t, x, y, z\}$ by the line element

$$ds^2 = g_{ab}dx^a dx^b = -\frac{dt^2}{U^2} + U^2(dx^2 + dy^2 + dz^2), \quad (4.1)$$

and electromagnetic one-form potential

$$A_a dx^a = \frac{1}{U} dt, \quad (4.2)$$

where $U(x, y, z)$ is any function of the *spatial* coordinates which satisfies Laplace’s equation $\nabla^2 U = 0$ on three-dimensional auxiliary Euclidean space. The linearity of Laplace’s equation means that solutions to (4.1) may be generated by linear superposition, a feature which is rare in general relativity due to the inherent non-linearity of Einstein’s field equations.

The Majumdar–Papapetrou solution to the Einstein–Maxwell equations was discovered independently by Majumdar [202] and Papapetrou [203] in 1947. The Majumdar–Papapetrou spacetime was studied by Hartle and Hawking [204], who found that for a particular solution U , the Majumdar–Papapetrou geometry describes a static configuration of N extremally charged Reissner–Nordström black holes.

4.2.1 Majumdar–Papapetrou di-hole

The Majumdar–Papapetrou solution that describes an assemblage of N extremal Reissner–Nordström black holes in static equilibrium is given by the line element (4.1) and electromagnetic one-form potential (4.2) with

$$U = 1 + \sum_{j=1}^N \frac{M_j}{\sqrt{(x-x_j)^2 + (y-y_j)^2 + (z-z_j)^2}}, \quad (4.3)$$

where M_j is the mass of the j th black hole, and its location in the background spatial coordinates is $(x, y, z) = (x_j, y_j, z_j)$. Each black hole is extremally charged ($Q_j = M_j$); the

gravitational attraction and electrostatic repulsion between the black holes are in perfect balance, ensuring a static spacetime.

We consider the Majumdar–Papapetrou binary black hole (or *di-hole*) solution, where (4.3) is given by

$$U = 1 + \frac{M_+}{\sqrt{x^2 + y^2 + (z - z_+)^2}} + \frac{M_-}{\sqrt{x^2 + y^2 + (z - z_-)^2}}. \quad (4.4)$$

The two black holes of mass M_{\pm} are located on the z -axis at $z_{\pm} = \pm \frac{dM_{\mp}}{M_+ + M_-}$, where d is the coordinate separation between the black holes in the background coordinates, and the centre of mass is at the origin of the coordinate system. The Majumdar–Papapetrou di-hole spacetime is symmetric about the axis connecting the black holes – here the z -axis. The black holes’ event horizons are located at the “points” $(x, y, z) = (0, 0, z_{\pm})$; these are in fact null surfaces of topology $\mathbb{R} \times S^2$, with non-zero area [204]. The geometry (4.1) with (4.4) is an example of a Weyl spacetime (see Section 2.3.3).

4.3 Null geodesic structure

To analyse the lensing effects of the Majumdar–Papapetrou di-hole, it is necessary to study the null geodesic structure of the spacetime. In this section, we briefly review the Lagrangian and Hamiltonian formalism for null geodesics, before analysing the null geodesic structure of the Majumdar–Papapetrou di-hole geometry. For a more detailed review of geodesic motion in general relativity, see Section 2.2.

4.3.1 Lagrangian formalism

As outlined in Section 2.2.2, a standard approach to solving for the geodesic motion on curved spacetime is to adopt the Lagrangian formalism [29]. In this scheme, one begins with the action, which is a functional of the spacetime coordinates, given by $S = \int L d\lambda$, with Lagrangian $L = \frac{1}{2}g_{ab}\dot{q}^a\dot{q}^b$. Here, $q^a(\lambda)$ is a spacetime path, and an overdot denotes the derivative with respect to the affine parameter λ . Along geodesics, the Lagrangian is conserved, with $L = 0$ in the case of null geodesics. The geodesics are given by solutions to the Euler–Lagrange equations,

$$\frac{d}{d\lambda} \left(\frac{\partial L}{\partial \dot{q}^a} \right) = \frac{\partial L}{\partial q^a}. \quad (4.5)$$

For null geodesics on the Majumdar–Papapetrou geometry (4.1), the Lagrangian takes the form

$$L = \frac{1}{2} \left[-\frac{1}{U^2} \dot{t}^2 + U^2 (\dot{x}^2 + \dot{y}^2 + \dot{z}^2) \right] = 0, \quad (4.6)$$

where U is given by (4.4) for the Majumdar–Papapetrou di-hole. Inserting the Lagrangian (4.6) into the Euler–Lagrange equations (4.5) yields a system of four second-order coupled non-linear ordinary differential equations for the path in spacetime $q^a(\lambda) = (t(\lambda), x(\lambda), y(\lambda), z(\lambda))$.

We recall that the Lagrangian (4.6) is not unique: any Lagrangian which yields the same equations of motion is valid (see Section 2.2.2). Some alternative choices of Lagrangian for the Majumdar–Papapetrou di-hole are outlined in [201].

4.3.2 Hamiltonian formalism

In the Hamiltonian approach, one uses the Lagrangian function to define the canonical momenta as $p_a = \frac{\partial L}{\partial \dot{q}^a} = g_{ab}\dot{q}^b$. Performing a Legendre transformation $H = \dot{q}^a p_a - L$, the corresponding Hamiltonian is then given in coordinates $\{q^a, p_b\}$ by the Hamiltonian function $H(q^a, p_a) = \frac{1}{2}g^{ab}p_a p_b$, where g^{ab} are the contravariant components of the metric tensor. Geodesics of the spacetime geometry are then the integral curves of Hamilton's equations, a system of first-order differential equations given by

$$\dot{q}^a = \frac{\partial H}{\partial p_a}, \quad \dot{p}_a = -\frac{\partial H}{\partial q^a}. \quad (4.7)$$

Along geodesics, the Hamiltonian is conserved, with $H = 0$ in the case of null rays. (See Section 2.2.3 for a review.)

In isotropic Cartesian coordinates $\{t, x, y, z\}$, the Hamiltonian for null geodesics on the Majumdar–Papapetrou di-hole spacetime takes the form

$$H = \frac{1}{2} \left[-U^2 p_t^2 + \frac{1}{U^2} (p_x^2 + p_y^2 + p_z^2) \right]. \quad (4.8)$$

Hamilton's equations (4.7) with Hamiltonian function (4.8) are given by

$$\dot{t} = -U^2 p_t, \quad \dot{p}_t = 0, \quad \dot{x} = \frac{p_x}{U^2}, \quad \dot{p}_x = -\frac{\partial U}{\partial x} \left[U p_t^2 + \frac{1}{U^3} (p_x^2 + p_y^2 + p_z^2) \right], \quad (4.9)$$

with similar expressions for y and z by symmetry of (4.8). The metric is independent of coordinate time t , so the momentum $p_t = -E$ is a constant of motion. Rescaling p_t is equivalent to rescaling the affine parameter λ ; we may therefore choose to fix $p_t = -1$ in the following analysis without loss of generality.

Rearranging the Hamiltonian constraint $H = 0$ for the Hamiltonian function (4.8), and inserting the result into the equations for \dot{p}_i given in (4.9), we see that the geodesic equations for the *spatial* components of the four-momentum can be expressed in the form $\dot{p}_i = - (U^2)_{,i}$. This resembles Newton's second law of motion, in which the force is expressed as the gradient of a scalar potential.

The diffeomorphism invariance of general relativity allows us to perform a change of coordinates without altering the physics of the problem at hand. We therefore choose to recast the problem in a coordinate system which is well-adapted to the symmetry of the configuration. The Majumdar–Papapetrou di-hole geometry is symmetric about the z -axis; it is therefore natural to choose cylindrical polar coordinates $\{t, \rho, \phi, z\}$, related to Cartesian coordinates in the standard way: $\rho^2 = x^2 + y^2$, $\phi = \arctan(\frac{y}{x})$. Having set $p_t = -1$, the Hamiltonian (4.8) now reads

$$H = \frac{1}{2} \left[\frac{1}{U^2} (p_\rho^2 + p_z^2 + \rho^2 p_\phi^2) - U^2 \right]. \quad (4.10)$$

In cylindrical coordinates, it is straightforward to see that ∂_t (∂_ϕ) is a Killing vector which generates time-translational (axial) symmetry. Hence, the coordinates t and ϕ are ignorable, and their conjugate momenta are conserved along geodesics. The physical interpretations of $-p_t$ and p_ϕ are, respectively, the total energy and axial angular momentum of the photon at infinity.

For the Majumdar–Papapetrou geometry in four spacetime dimensions, the full phase space is eight-dimensional, spanned by the four spacetime coordinates and their conjugate momenta $\{q^a, p_a\}$. The conserved momenta allow us to focus on a reduced four-dimensional phase space on which the dynamics is governed by two pairs of conjugate variables $\{\rho, z, p_\rho, p_z\}$, and the null condition $H = 0$. This constraint allows us to express one of the phase space coordinates (e.g. p_z) in terms of the other three.

Recall from Section 2.2.1 that the null geodesics of the conformally related metrics g_{ab} and $\tilde{g}_{ab} = \Omega^2 g_{ab}$ coincide, where $\Omega > 0$ is a (smooth) function of the spacetime coordinates. Performing a conformal transformation with $\Omega = \frac{1}{U}$ allows us to recast the Hamiltonian (4.10) in canonical form as

$$\tilde{H} = \frac{1}{2} (p_\rho^2 + p_z^2) + V, \quad V(\rho, z) = -\frac{1}{2\rho^2} (h - p_\phi)(h + p_\phi), \quad (4.11)$$

where we have factorised the potential $V(\rho, z)$ by introducing the *effective potential* (or *height function*)

$$h(\rho, z) = \rho U^2. \quad (4.12)$$

The Hamiltonian (4.11) will be used throughout our analysis of null geodesics on the Majumdar–Papapetrou di-hole spacetime, with the effective potential (4.12) playing a key role. (We hereafter omit the tilde from the Hamiltonian (4.11), simply writing H .) Using the Hamiltonian (4.11), the geodesic equations (i.e., Hamilton’s equations) are

$$\dot{\rho} = p_\rho, \quad \dot{z} = p_z, \quad \dot{\phi} = \frac{p_\phi}{\rho^2}, \quad (4.13)$$

$$\dot{p}_\rho = U^2 \frac{\partial(U^2)}{\partial \rho} + \frac{p_\phi^2}{\rho^3}, \quad \dot{p}_z = U^2 \frac{\partial(U^2)}{\partial z}, \quad \dot{p}_\phi = 0. \quad (4.14)$$

4.3.3 Integrability and chaotic scattering

The Newtonian analogue of geodesic motion on the Majumdar–Papapetrou spacetime is the problem of two fixed centres (also known as Euler’s three-body problem or the two-centre Kepler problem) [142, 219]. This is a special case of the classical three-body problem of three point masses described by Newton’s laws of motion and Newton’s law of universal gravitation. The general three-body problem has no closed-form solution.

Remarkably, Euler’s restricted three-body problem is soluble. In the case of two fixed centres, the equations governing the motion of a test particle are *separable* in spheroidal coordinates. This is due to the fact that, besides the total energy and one component of the angular momentum, there exists an additional “hidden” constant of motion – Whittaker’s constant [220, 221] (cf. Carter’s constant [47] for geodesic motion on Kerr spacetime). The problem is therefore Liouville integrable; see Section 2.6.

In contrast to motion in Euler’s three-body problem, geodesic motion on the Majumdar–Papapetrou di-hole spacetime – as well as other relativistic generalisations, such as the Weyl–Bach di-hole [222, 223] – is not integrable, as there is no known analogue of Whittaker’s constant [224].¹ The lack of Liouville integrability on the Majumdar–Papapetrou

¹In [223], the authors demonstrate that, for the special case of equal-mass black holes, the \mathbb{Z}_2 -symmetric Weyl–Bach di-hole admits a $(2+1)$ -dimensional totally geodesic submanifold – the equatorial plane – on which the geodesic motion is Liouville integrable. This is also explored by Assumpção *et al.* [225].

di-hole spacetime results in rich phenomena. Contopoulos revealed that the Majumdar–Papapetrou di-hole exhibits chaotic dynamics and self-similarity [140–142]. This paved the way for deeper insights into chaos in binary systems in general relativity from Contopoulos and collaborators [205–208]; Yurtsever [209]; Cornish, Dettmann, Frankel, Gibbons, Levin and collaborators [134, 210–215]; and others [127, 138, 139, 216–218]. In this chapter, we extend on these influential perspectives, advancing the view that the scattering of photons by a pair of black holes in the Majumdar–Papapetrou di-hole spacetime is a natural example of *chaotic scattering* (see Section 2.6.5 for a review).

For the Majumdar–Papapetrou di-hole, a scattering problem is set up as follows. We consider the evolution of null geodesics (light rays) for different initial conditions (impact parameters). The final state of the photon is characterised by three possible outcomes. A photon may (i) cross the event horizon of the upper black hole; (ii) cross the event horizon of the lower black hole; (iii) escape to spatial infinity. In this system, the scattering singularities correspond to the null geodesics which are non-escaping (“perpetual”) orbits in forward time ($\lambda \rightarrow \infty$). The non-escaping orbits are neither scattered nor absorbed by the black holes. The scattering singularities, which comprise the *repellor* (Section 2.6.5), can be divided into two classes: the subset of *periodic* orbits; and the subset of *aperiodic* orbits. In the following sections, we shall explore the structure of these subsets of the initial data, and analyse some examples of non-escaping orbits.

4.4 Planar geodesics and one-dimensional shadows

In this section, we consider null geodesics on the Majumdar–Papapetrou di-hole spacetime in standard coordinates $\{t, x, y, z\}$, using the Hamiltonian formalism outlined in Section 4.3.2. The line element is given by (4.1) and the electromagnetic one-form potential by (4.2), where the function U appearing in both quantities is given by (4.4). For simplicity, we consider the symmetric case of equal-mass black holes ($M_+ = M_- = M$) located at the points $z = \pm \frac{d}{2}$, where d is the coordinate separation between the black holes. We employ units in which $M = 1$.

We shall examine a highly symmetric scattering problem, in which null rays are confined to the meridian plane ($y = 0$). The conservation of $p_\phi = xp_y - yp_x$ ensures that rays which start in the meridian plane (with $y = 0 = p_y$) remain in the plane. We consider rays which start at the centre of mass $(x, z) = (0, 0)$ between the black holes. The scattering scenario is one-dimensional: the initial state of the system is characterised by the angle α made by the momentum two-vector (p_x, p_z) and the positive x -axis. The final state of the scattering process is described by three distinct possibilities: (i) the photon crosses the event horizon of the upper black hole; (ii) the photon crosses the event horizon of the lower black hole; (iii) the photon escapes to spatial infinity.

The initial data (with $\lambda = 0$) for this scattering problem are taken to be

$$t = x = y = z = 0, \quad p_t = -1 \quad p_y = 0, \quad p_x = U_0^2 \cos \alpha, \quad p_z = U_0^2 \sin \alpha, \quad (4.15)$$

where α is the initial angle and $U_0 = U(0, 0, 0) = 1 + \frac{4}{d}$. Taking $d = 2$ as a default value, we numerically evolve the null geodesic equations (Hamilton’s equations) with initial data

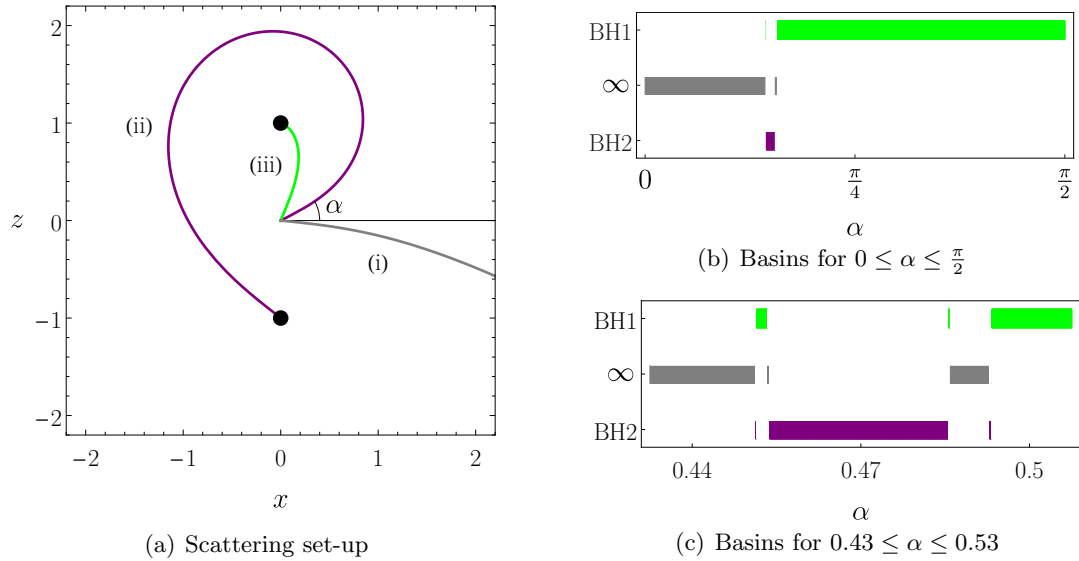


Figure 4.1. Set-up of the one-parameter scattering problem and corresponding one-dimensional exit basins. (a) Rays which emanate from the centre of mass between the black holes with initial scattering angle α may: (i) escape to spatial infinity; (ii) fall into the lower black hole; (iii) fall into the upper black hole. (b) One-dimensional exit basin diagram for $\alpha \in [0, \frac{\pi}{2}]$. (c) Magnification of the exit basin diagram from (b), with $\alpha \in [0.43, 0.53]$.

given by (4.15) for $0 \leq \alpha \leq \frac{\pi}{2}$; we need only consider rays whose initial momentum two-vector (p_x, p_z) points into the first quadrant ($x \geq 0, z \geq 0$) due to the symmetry of the scattering problem.

Figure 4.1(a) shows examples of null geodesics emanating from the origin for different values of the initial angle α . We see that, for $\alpha \sim 0$, the photon will escape to spatial infinity; if $\alpha \sim \frac{\pi}{2}$, the ray will fall directly into the upper black hole; for intermediate values of α , the ray may orbit the upper black hole before plunging into the lower black hole, or it may pass between the black holes before escaping to spatial infinity. Figure 4.1(b) shows the final fate of the photon as a function of α . This is an example of a one-dimensional *exit basin diagram* (see Section 2.6.5), which allows us to gain insight into the nature of scattering by a pair of fixed black holes.

The exit basin diagram presented in Figure 4.1(b) can be thought of as a one-dimensional binary black hole shadow for an observer located at the centre of mass. For the one-parameter scattering problem outlined above, the shadow (or exit basin) of the upper/lower black hole can be defined formally as

$$\mathcal{B}_{\pm} = \lim_{\varepsilon \rightarrow 0} \left\{ \alpha \in [0, 2\pi) \mid x^2(\lambda; \alpha) + (z(\lambda; \alpha) - z_{\pm})^2 \leq \varepsilon, \lambda \in (-\infty, 0) \right\}, \quad (4.16)$$

where $(x, z) = (0, z_{\pm})$ is the location of the upper/lower black hole. Here, the notation $x(\lambda; \alpha)$ means $x(\lambda)$ with initial data α specified by (4.15). The binary black hole shadow is then the union of the two exit basins: $\mathcal{B}_S = \mathcal{B}_+ \cup \mathcal{B}_-$. Similarly, the basin corresponding to spatial infinity can be defined as

$$\mathcal{B}_{\infty} = \lim_{R \rightarrow \infty} \left\{ \alpha \in [0, 2\pi) \mid \lim_{\lambda \rightarrow -\infty} r(\lambda; \alpha) \geq R \right\}, \quad (4.17)$$

where $r = \sqrt{x^2 + z^2}$ is the radial distance from the origin in the meridian plane. The set of scattering singularities in the initial data is given by $[0, 2\pi) \setminus (\mathcal{B}_S \cup \mathcal{B}_\infty)$.

In Figure 4.1(c), we magnify an interesting region of the exit basin diagram from Figure 4.1(b) which demonstrates rich structure for intermediate scattering angles $\alpha \in [0.43, 0.53]$. Inspecting this figure, it appears that the one-dimensional shadows exhibit intricate detail on smaller and smaller scales, hinting at self-similarity. In Section 4.5, we will analyse the trajectories of the planar one-parameter scattering problem in more detail. Then, in Section 4.6.1, we will then demonstrate that this analysis can be used to understand the rich structure hinted at in Figure 4.1(c) in greater detail.

4.5 Symbolic dynamics

4.5.1 Chaotic scattering and symbolic dynamics

One way to demonstrate that a system is chaotic is through the use of *symbolic dynamics*, which provides a succinct description of the dynamics of chaotic systems. Symbolic dynamics describes the topology of trajectories in phase space, encoding trajectories as a string of abstract symbols. This method provides a coordinate-invariant method of characterising chaos, which is of particular importance when studying the dynamics of general-relativistic systems, due to the general covariance of the theory. Furthermore, symbolic dynamics can be studied *analytically*, despite the fact that the equations of motion are themselves non-integrable.

In the study of chaotic scattering, the subset of the initial data corresponding to trajectories which asymptote to unstable *periodic* orbits, denoted by Ω_P , is particularly important. Moreover, the subset of unstable *aperiodic* orbits Ω_A also plays an important role. The development of a suitable symbolic code allows one to develop a greater understanding of the rich structure of the strange repeller $\Omega_R = \Omega_P \cup \Omega_A$ (see Section 2.6.5).

The aim of this section is to develop a symbolic dynamics for the Majumdar–Papapetrou di-hole. Rather than give a formal overview of symbolic dynamics, we will instead review two symbolic codes for the Gaspard–Rice three-disc model [120]. The first symbolic coding, referred to here as *collision dynamics* was introduced in [119]; the second, which we refer to as *decision dynamics* is an alternative symbolic code which allows a greater understanding of the ordering of non-escaping orbits in the initial data, i.e., the structure of the repeller Ω_R . We will then turn our attention to the development of collision dynamics and decision dynamics for the Majumdar–Papapetrou di-hole.

4.5.2 Symbolic dynamics for the Gaspard–Rice system

The canonical example of a chaotic scatterer is the three-disc model proposed by Eckhardt [119], which has been studied extensively by Gaspard and Rice [120], and others [226, 227]. In this two-dimensional system, an incoming particle, whose initial state is characterised by a pair of initial conditions (i.e., position and momentum), undergoes perfectly elastic collisions with three fixed hard discs (typically of equal radius and situated at the vertices

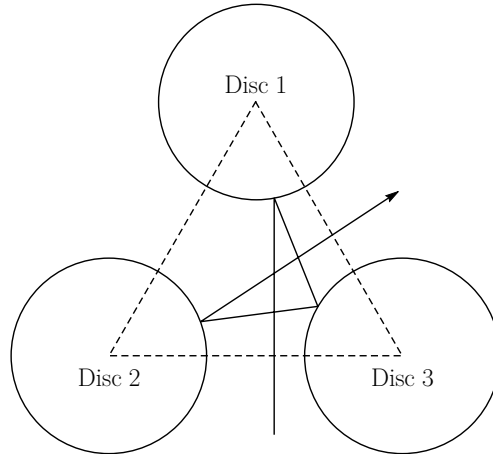


Figure 4.2. Scattering geometry for the Gaspard–Rice three-disc model. A particle undergoes perfectly elastic scattering with three hard discs, centred on the vertices of an equilateral triangle. Using collision dynamics, the reference trajectory – which first hits Disc 1, then Disc 3, then Disc 2, before escaping to infinity – can be encoded as the sequence of symbols 132.

of an equilateral triangle) until it escapes to infinity. The set-up is shown in Figure 4.2.

The qualitative features of the system can be understood through the use of symbolic dynamics. The standard approach is to assign a label from the alphabet $\mathcal{A} = \{1, 2, 3\}$ to each of the discs. A trajectory is then labelled by a sequence of digits from this alphabet, each of which records the collision of the particle with the disc of the same label. One must also impose the requirement that no digit follows itself (i.e., the particle may not collide with the same disc twice in a row). For example, the sequence 132 labels a trajectory which first hits Disc 1, then Disc 3, then Disc 2, before escaping from the scattering region (see Figure 4.2).

In principle, there exist sequences of arbitrary length in this symbolic dynamics representation. The set of perpetual non-escaping orbits are encoded by (bi-)infinite sequences of digits from \mathcal{A} . Formally, these are sequences \mathcal{X} which belong to the set $\Sigma = \{\mathcal{X} = (\mathcal{X}_i)_{i \in \mathbb{F}} \mid \mathcal{X}_i \in \mathcal{A}\}$, where $\mathbb{F} = \mathbb{N}$ ($\mathbb{F} = \mathbb{Z}$) for infinite (bi-infinite) sequences. Here, we are mostly concerned with trajectories which start far from the scattering region. We therefore only consider finite sequences (escaping trajectories) or infinite sequences (non-escaping trajectories), rather than bi-infinite sequences. An important subset of Σ is the *countable* set of (eventually) recurring sequences describing (asymptotically) periodic orbits; the embedding of this set in Σ is akin to the embedding of \mathbb{Q} in \mathbb{R} . There also exists an *uncountable* set of non-recurring infinite sequences, which describe aperiodic orbits.

There is an alternative approach to labelling the trajectories of the three-disc system. After each collision with a disc, a trajectory can be continued in three possible ways: the particle may hit the disc on the left of the current disc (0); it may collide with the disc on the right (2); or it may escape the scattering region (1). In this scheme, the perpetual orbits are described by (bi-)infinite sequences which do not contain the digit 1. There is a natural mapping between the symbolic representation of non-escaping orbits in this picture and the ternary representation of the middle-thirds Cantor set; this will be discussed in

more detail in Section 4.6.2.

We distinguish between these two types of symbolic dynamics as follows: the latter approach is referred to as *decision dynamics*, whilst the former is termed *collision dynamics*. In decision dynamics, repeated neighbouring digits are permitted; this is not the case in collision dynamics. Moreover, the representation of orbits in decision dynamics provides natural insight into the structure of the strange repeller Ω_R – the subset of initial data corresponding to the non-escaping perpetual orbits.

The decision dynamics representation using the alphabet $\{0, 1, 2\}$ is useful for understanding the structure of Ω_R . However, in open Hamiltonian systems, one is typically concerned with the *escaping* trajectories, as well as the non-escaping ones. In the three-disc model, one may wish to record the exit through which the particle left the scattering region. We may therefore reformulate our symbolic code to account for this as follows. A particle which has just collided with a disc may: escape through the exit on the left (0); collide with the disc on the left (1); escape through the exit opposite (2); collide with the disc on the right (3); escape through the exit on the right (4).

In this symbolic code, one can describe non-escaping orbits as infinite sequences which do not contain the digits 0, 2, or 4. A key point is that finite sequences encode *families* of trajectories, which correspond to open sets in the initial data. Conversely, infinite sequences describe *unique* non-escaping orbits, which correspond to scattering singularities in the initial data.

4.5.3 Symbolic dynamics for the Majumdar–Papapetrou di-hole

We now turn our attention to a symbolic dynamics representation of null geodesics on the Majumdar–Papapetrou di-hole spacetime. We develop the symbolic coding by considering a null geodesic in a congruence which has reached a “decision point”. This is illustrated in Figure 4.3. The geodesic may follow a path around the opposite black hole in the same sense (0); around the other black hole in the opposite sense (2); or around the same black hole in the same sense (4). Moreover, the photon could fall into a black hole (1); or escape to infinity (3). In order to avoid double-counting, we do not enumerate the possibility of falling into the “other” black hole; this is accounted for at a previous or subsequent decision point. The rays which plunge directly into a black hole, or directly escape to infinity do not generate interesting structure in the exit basins (shadows).

Null rays confined to the meridian plane in the Majumdar–Papapetrou di-hole spacetime can be described using the base-5 symbolic alphabet $\mathcal{A} = \{0, 1, 2, 3, 4\}$. The non-escaping orbits – which remain in the strong-field region without plunging into a black hole or escaping to infinity – are described by infinite sequences of *even* digits from the alphabet \mathcal{A} , i.e., sequences from the set

$$\Sigma_R = \{\mathcal{X} = (\mathcal{X}_i)_{i \in \mathbb{N}} \mid \mathcal{X}_i \in \{0, 2, 4\}\}, \quad (4.18)$$

The rays which escape the system (by crossing an event horizon or reaching spatial infinity) are represented by sequences of finite length $k \in \mathbb{N}$, whose first $k - 1$ digits are even, and which terminate in the digit 1 (absorbed) or 3 (scattered); formally, these are sequences

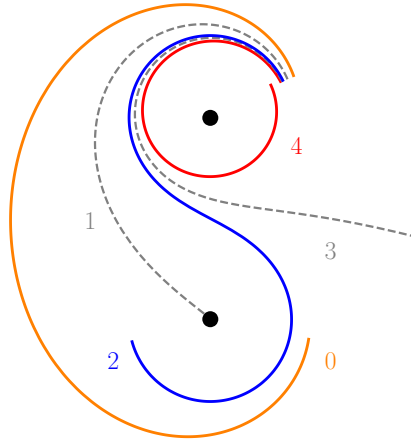


Figure 4.3. Schematic diagram of “decision dynamics” – a version of symbolic dynamics – for the Majumdar–Papapetrou di-hole.

from the set

$$\Sigma_E = \left\{ \mathcal{X} = (\mathcal{X}_i)_{i=1}^k \mid \mathcal{X}_i \in \{0, 2, 4\}, 1 \leq i \leq k-1, \mathcal{X}_k \in \{1, 3\}, \forall k \in \mathbb{N} \right\}. \quad (4.19)$$

Sequences in Σ_E (which are of finite length) describe families of orbits, which are given by open sets in the initial data. (In fact, any sequence of finite length describes a family of orbits; such a sequence need not belong to Σ_E . For example, the sequence 024 represents the family of trajectories which first make decision 0, then decision 2, then decision 4. The corresponding initial conditions form a set of finite measure.) Sequences which belong to Σ_R encode unique perpetual orbits, each of which generates a scattering singularity in the initial data. These constitute the repeller Ω_R .

An important subset of the perpetual orbits is the set of periodic orbits. In the decision dynamics scheme, periodic orbits are represented by repeating sequences from Σ_R . For example, the recurring sequence $\bar{0} = 000 \dots$ corresponds to the periodic orbit which loops around both black holes; the sequence $\bar{2}$ corresponds to a figure-of-eight orbit; and $\bar{4}$ describes the null geodesic which orbits an individual black hole. These three periodic orbits, and an additional example, are shown in Figure 4.4. A generic non-escaping orbit need not be periodic, however. The representation of aperiodic orbits is given by non-recurring sequences from Σ_R .

The encoding of orbits using a base-5 symbolic alphabet provides a natural mapping between the initial data corresponding to the perpetual orbits and the elements of the 5-adic Cantor set. (This will be discussed in more detail in Section 4.6.)

Cornish and Gibbons [213] provide an alternative symbolic coding for planar unstable photon orbits in the Majumdar–Papapetrou di-hole spacetime.² In their approach, a null

²In fact, Cornish and Gibbons [213] analyse a one-parameter family of solutions to the Einstein–Maxwell–dilaton theories, of which the Majumdar–Papapetrou di-hole geometries are a subfamily. The authors explore transitions between regular and chaotic motion as the dilaton coupling is varied, and the integrable Kaluza–Klein limit. The article also contains a discussion of the topological entropy of the symbolic coding.

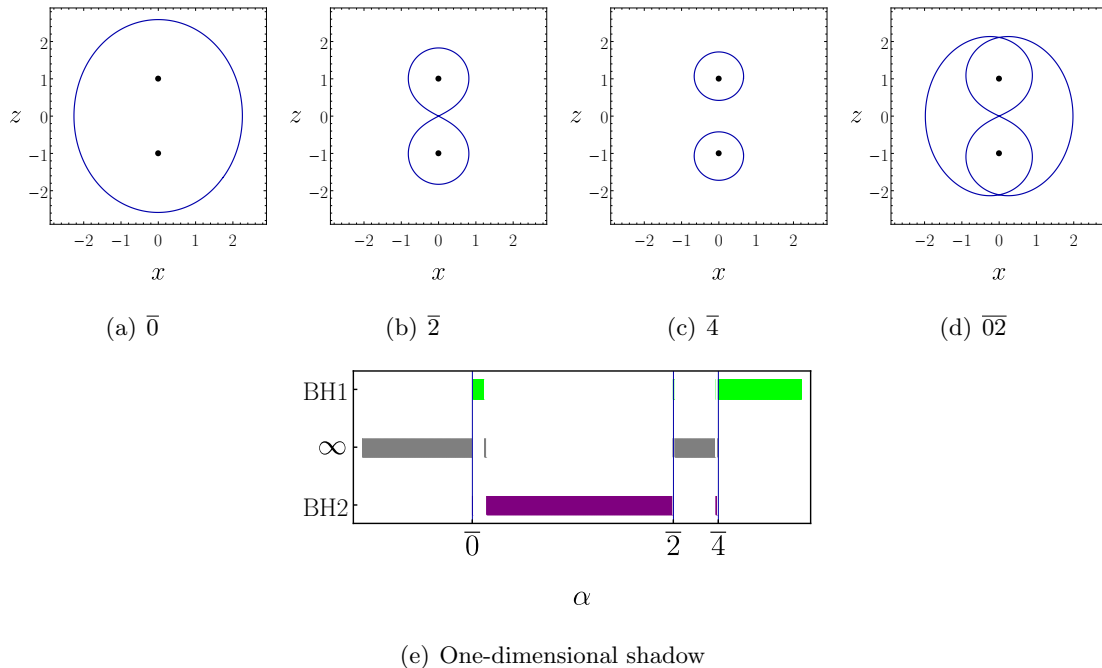


Figure 4.4. Examples of planar periodic null orbits with zero angular momentum p_ϕ . We choose $M_\pm = 1$ and $d = 2$ as default parameter values. In our symbolic code, these orbits are represented by the recurring sequences (a) $\bar{0}$, (b) $\bar{2}$, (c) $\bar{0}$, and (d) $\bar{02}$. The orbits (a)–(c) are referred to as the *fundamental* periodic orbits. Orbit (d) undergoes “dynamical transitions” between orbits (a) and (b); this is reflected in its symbolic representation. (e) One-dimensional binary black hole shadow. The vertical lines show the location of the fundamental orbits (a)–(c) in the initial data.

geodesic is described by a sequence of digits from the alphabet $\{+1, 0, -1\}$ which record the passage of the geodesic through three “windows” placed on the symmetry axis. The digits $+1$, 0 , and -1 correspond, respectively, to the intersection of the geodesic with the open intervals $z \in (z_+, +\infty)$, $z \in (z_-, z_+)$, and $z \in (-\infty, z_-)$. We classify this symbolic code as an example of collision dynamics (cf. the three-disc model in Section 4.5.2). In this work, we favour decision dynamics for the Majumdar–Papapetrou di-hole, as it provides a better description of the structure of the strange repeller Ω_R , and it may be used to describe both the escaping and non-escaping trajectories.

Alonso *et al.* [218] employ a yet another symbolic code to describe non-escaping orbits in the meridian plane. The authors use a type of collision dynamics, describing orbits with symbols from the alphabet $\{+, o, -\}$, depending on where the ray intersects the symmetry axis. This choice of symbolic coding is similar to that of Cornish and Gibbons [213].

4.6 Cantor-like structure of the one-dimensional shadow

4.6.1 Ordering of non-escaping orbits in the initial data

The non-escaping orbits are of importance for a number of reasons. Firstly, in a general open Hamiltonian dynamical system, the non-escaping trajectories correspond to scattering singularities in the initial data: they form the repeller Ω_R (see Section 2.6.5). The

structure of the strange repeller provides useful information about the dynamics of chaotic scattering systems. Secondly, in the context of scattering and absorption of null geodesics by black holes, the perpetual orbits lie on the boundary between the absorbed and scattered rays; that is, they correspond to the *boundary* of the black hole shadow. In order to understand the structure of the shadows of a pair of fixed black holes, one must therefore understand the structure of the perpetual orbits in the associated scattering problem.

We now apply the symbolic code developed in Section 4.5.3 to gain insight into the ordering and organisation of the non-escaping orbits in the initial data, for the one-dimensional scattering problem described in Section 4.4. Consider the one-dimensional shadow (or exit basin diagram) for $\alpha \in [0, \frac{\pi}{2}]$, presented in Figure 4.1(b). Note that we need only consider this restricted domain due to the symmetry of the equal-mass scattering problem. As described in Section 4.4, all rays with $\alpha \sim 0$ will escape to infinity, and rays with $\alpha \sim \frac{\pi}{2}$ will plunge directly into the upper black hole. For intermediate values of α , we observe interesting structure in the exit basins. In particular, we highlight three regions of interest; see Figure 4.7(a). Outside of these open intervals, there are no scattering singularities: all trajectories are either absorbed or scattered by the system. However, each of these open intervals contains infinitely many scattering singularities.

To the left of the left-hand interval, all trajectories escape to infinity. Between the left and middle intervals, all rays cross the event horizon of the lower black hole. All initial conditions which lie between the middle and right-hand intervals correspond to rays which escape to infinity. Finally, to the right of the right-hand interval, all trajectories fall into the upper black hole.

Let us consider what happens as one increases the scattering parameter α continuously, starting from $\alpha = 0$. It is clear from the exit basin diagram that, for α sufficiently small, all rays escape to infinity. This is the case until we reach some critical value of α which corresponds to the *first* perpetual orbit. This orbit asymptotes to the perpetual orbit which has symbolic representation $\bar{0}$ in decision dynamics; see Figure 4.4(a).

In a similar manner, let us consider the effect of continuously decreasing α , starting from $\alpha = \frac{\pi}{2}$. Initially, all trajectories fall into the upper black hole, until we reach a critical value of α , corresponding to the *last* perpetual orbit (i.e., the one with the largest value of $\alpha \in [0, \frac{\pi}{2}]$). This is a periodic orbit, which loops only around the upper black hole. In decision dynamics, this orbit is represented by the recurring sequence $\bar{4}$; see Figure 4.4(c).

At first sight, it appears from Figure 4.4(e) that the ordering of the perpetual orbits in the initial data exactly matches the natural ordering of their symbolic representations, when the latter are interpreted as elements of the unit interval in base-5; however, this is not the case.

The set of all possible symbolic sequences for planar trajectories is $\Sigma = \Sigma_R \cup \Sigma_E$, where sequences in Σ_R describe non-escaping trajectories, and sequences in Σ_E describe trajectories which fall into a black hole or escape to infinity; see Section 4.5.3. We now define an *ordering function* $F: \Sigma \rightarrow [0, 1]$, which maps a decision dynamics sequence to an element of the unit interval, represented in base-5. We demand that the ordering of elements in $\text{im}(F) \subseteq [0, 1]$ exactly matches the ordering of the scattering singularities in the initial data; that is, $F(\mathcal{X}^{(1)}) < F(\mathcal{X}^{(2)})$ if and only if $\alpha^{(1)} < \alpha^{(2)}$, where $\mathcal{X}^{(i)} \in \Sigma$ is

any sequence and $\alpha^{(i)}$ is its corresponding initial condition.

The ordering function is defined as follows. We first define a *parity-reordering operation* $\mathcal{P}: \Sigma \rightarrow \Sigma$, which maps each digit \mathcal{X}_i of a sequence $\mathcal{X} = \mathcal{X}_1\mathcal{X}_2\mathcal{X}_3 \cdots \in \Sigma$ to a digit $\tilde{\mathcal{X}}_i$ of the sequence $\mathcal{P}(\mathcal{X})$, according to the following procedure which keeps track of parity P . We begin with $P = +1$ and consider the sequence \mathcal{X} from left to right, examining each digit \mathcal{X}_i in turn, as follows.

(1) Set

$$\tilde{\mathcal{X}}_i = \begin{cases} \mathcal{X}_i, & \text{if } P = +1; \\ 4 - \mathcal{X}_i, & \text{if } P = -1. \end{cases} \quad (4.20)$$

(2) Reverse the parity ($P \mapsto -P$) if $\mathcal{X}_i = 2$.

(3) Iterate $i \mapsto i + 1$.

We note that $\mathcal{P}^2(\mathcal{X}) = \mathcal{X}$, so \mathcal{P} is an involution. The parity-reordering operation can be understood in terms of the geometry of rays: the digit 2 describes a null geodesic which passes between the black holes, changing the sense of the orbit from clockwise to anti-clockwise and vice versa. Hence, the ordering of digits in our symbolic code must be reversed. (It will be convenient to write $\mathcal{P}(\mathcal{X}) = \tilde{\mathcal{X}}$ to denote parity-reordered sequences.)

Next, we define $f: \Sigma \rightarrow [0, 1]$ to be the function which maps a (parity-reordered) sequence to a real number in the unit interval represented in base-5. We take $f(\mathcal{Y}) = 0.\mathcal{Y}_1\mathcal{Y}_2\mathcal{Y}_3 \cdots$, where $\mathcal{Y} = \mathcal{Y}_1\mathcal{Y}_2\mathcal{Y}_3 \cdots \in \Sigma$ is any (parity-reordered) sequence. Finally, we take the ordering function to be $F(\mathcal{X}) = (f \circ \mathcal{P})(\mathcal{X})$ for all sequences $\mathcal{X} \in \Sigma$.

To illustrate how F acts on sequences in Σ , consider the pair of recurring sequences $\mathcal{X}^{(1)} = \overline{20}$, $\mathcal{X}^{(2)} = \overline{24}$. The parity-reordered sequences are $\mathcal{P}(\mathcal{X}^{(1)}) = \overline{2420}$ and $\mathcal{P}(\mathcal{X}^{(2)}) = \overline{2024}$. We therefore find $F(\mathcal{X}^{(1)}) = 0.\overline{2420}$ and $F(\mathcal{X}^{(2)}) = 0.\overline{2024}$. Since $F(\mathcal{X}^{(2)}) < F(\mathcal{X}^{(1)})$, the perpetual orbit described by the sequence $\mathcal{X}^{(2)}$ will precede that represented by $\mathcal{X}^{(1)}$ in the initial data; in other words, $\alpha^{(2)} < \alpha^{(1)}$.

4.6.2 Constructing a Cantor-like set on the initial data

Cantor sets

The *middle-thirds Cantor set* \mathcal{C} is a paradigm of fractal geometry; its complex structure can be arrived at using a remarkably simple iterative procedure [114]. Starting with the closed unit interval $I = [0, 1]$, one first removes the open middle-third interval $(\frac{1}{3}, \frac{2}{3})$, which leaves two closed intervals $[0, \frac{1}{3}]$ and $[\frac{2}{3}, 1]$. At the next step, one removes the open middle-third of the remaining closed intervals. This is repeated *ad infinitum*, until one is left with a set of distinct points of zero Lebesgue measure which were not removed at any step.

In general, a set Λ is a *Cantor set* if it is a closed, totally disconnected, and perfect subset of I [95]. Recall that a set is *closed* if it contains all of its boundary points; a set is *totally disconnected* if it contains no intervals; and a set is *perfect* if every point in it is a limit point of other points in the set. Let A be a subset of a topological space X . A point

$p \in X$ is a *limit point* of A if every open neighbourhood of p contains at least one point of A which is different from p .

Related Cantor-like sets may be generated by removing multiple intervals at each step, or by varying the proportionate width of the interval(s) removed. A straightforward generalisation of the middle-thirds Cantor set is the *middle- γ Cantor set* (or *γ -Cantor set*), denoted \mathcal{C}_γ [95, 113, 114]. This can again be constructed iteratively on the unit interval: at each stage one removes the open middle interval of length $\gamma \in (0, 1)$. Clearly the case $\gamma = \frac{1}{3}$ gives rise to the standard middle-thirds Cantor set \mathcal{C} .

The middle-thirds Cantor set is an example of a *fractal* (see Section 2.6.5), a set which is self-similar under magnification [95]. Consider only those points of the Cantor set which lie in the left-hand interval $[0, \frac{1}{3}]$. Magnifying this region by a factor of three, the portion of the Cantor set in $[0, \frac{1}{3}]$ looks exactly the same as the original set in I . (More formally, the linear map $L(x) = 3x$ is a homeomorphism from the piece of the Cantor set in $[0, \frac{1}{3}]$ to the whole Cantor set.) This self-similarity is not restricted to one level of magnification: enlarging any part of the Cantor set at the k th iteration in its construction by a factor of 3^k yields the original set. The γ -Cantor set also exhibits self-similarity.

The fractal structure of the Cantor set (and its complement in the unit interval) can be understood in more depth by considering the ternary representation of numbers in the unit interval I . It is well-known that the middle-thirds Cantor set consists of all elements of I which can be represented in base-3 using only the digits 0 and 2; see Theorem 4.2 of [97], for example. The action of deleting the middle-third interval at the k th iteration is equivalent to removing from I all elements with a ternary representation whose k th digit is 1, and whose first $k - 1$ digits consist only of the digits 0 and 2.

Cantor basins

Consider the iterative process used to generate the γ -Cantor set. One can construct a pair of basins $\{B_i \mid i \in \{1, 2\}\}$, which we refer to as *Cantor basins*, from the complement of \mathcal{C}_γ in I as follows. Denote by J_k the union of the 2^{k-1} open intervals which are *removed* from I at the k th iteration in the construction of \mathcal{C}_γ . Now, add the open set J_k to the basin B_i , where $i = 1$ ($i = 2$) if k is odd (even). The basins can then be expressed as an infinite union of disjoint open sets:

$$B_1 = \bigcup_{l=1}^{\infty} J_{2l-1}, \quad B_2 = \bigcup_{l=1}^{\infty} J_{2l}. \quad (4.21)$$

Clearly, $\partial B_1 = \partial B_2 = \mathcal{C}_\gamma$, i.e., the basins B_i share a common boundary, namely the γ -Cantor set.

It is possible to generalise the procedure outlined above in order to construct an N -basin set $\{B_i \mid i \in \{1, 2, \dots, N\}, N \in \mathbb{N}\}$, where $\partial B_1 = \partial B_2 = \dots = \partial B_N = \mathcal{C}_\gamma$. Again, we denote by J_k the union of the open intervals removed at the k th step, and we define the N disjoint basins as

$$B_i = \bigcup_{k=i \bmod N} J_k, \quad i \in \{1, 2, \dots, N\}. \quad (4.22)$$

The construction of the Cantor basins associated with the middle-thirds Cantor set can be understood by considering the ternary representation of elements in the unit interval,

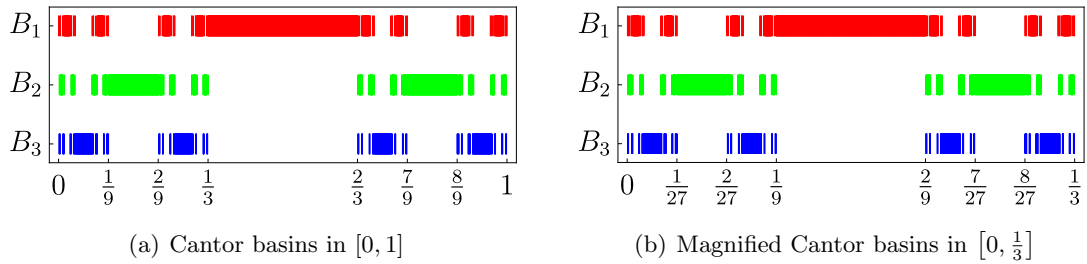


Figure 4.5. Three-colour Cantor basins constructed using the iterative procedure used to generate the middle-thirds Cantor set. (a) Three basins $\{B_1, B_2, B_3\}$ on the unit interval, coloured red, green and blue, respectively. The boundary of the basins is precisely the Cantor set. (b) Magnification of the basins of (a), showing the interval $[0, \frac{1}{3}]$. One can clearly see that the basins are self-similar under this magnification.

which can be viewed as a “symbolic dynamics” for the Cantor set; see Figure 4.1 of [97]. For simplicity, consider the case of two basins $\{B_1, B_2\}$. All points $p \in I$ with a ternary expansion whose first 1 appears in the k th slot will belong to B_1 (B_2) if k is odd (even). For example, all points with ternary expansion $0.021\dots$ are removed at the third iteration, so they will belong to B_1 (in the case of two basins). The boundary points (contained in neither basin) are numbers in I whose ternary expansion does not contain the digit 1. (These are precisely the elements of the Cantor set.) In the language of dynamical systems, the Cantor set (i.e., the boundary of the Cantor basins) is a strange repeller.

In the case of Cantor basins with $\gamma \neq \frac{1}{3}$, the symbolic alphabet $\{0, 1, 2\}$ does not correspond to the digits in the ternary representation of element of the unit interval; rather, the digits 0, 1 and 2 represent the left, middle, and right intervals, respectively, at each step in the construction of the set. For example, the symbolic code 021 (represented in base-3 by the number 0.201) corresponds to the interval which is arrived at by choosing first the closed left-hand interval (0); choosing the closed right-hand sub-interval at the next iteration (2); and finally choosing the open middle sub-interval (1). It is straightforward to see that, in the case of two basins $\{B_1, B_2\}$, the interval represented by the symbolic code 021 would belong to basin B_1 . We emphasise that it is *not* the case that each element of the open interval represented by 201 will have a ternary expansion which begins 0.201. This is only true for $\gamma = \frac{1}{3}$.

In Figure 4.5, we plot $N = 3$ Cantor basins which are generated from the construction of the canonical middle-thirds Cantor set \mathcal{C} on the unit interval $[0, 1]$. The basins B_1 , B_2 and B_3 are coloured red, green and blue, respectively. From Figure 4.5(a), one can see that the open middle-third interval $(\frac{1}{3}, \frac{2}{3})$, which is removed at the first step in the construction of the middle-thirds Cantor set, belongs to be basin B_1 ; the union of the open middle-third intervals of the two remaining intervals, given by $(\frac{1}{9}, \frac{2}{9}) \cup (\frac{7}{9}, \frac{8}{9})$, belong to the basin B_2 , and so on. The boundary of the three basins is the Cantor set. In Figure 4.5(b), we magnify the Cantor basins of Figure 4.5(a) by a factor of 3, considering the left-hand interval $[0, \frac{1}{3}]$. In this magnified figure, the self-similarity of the Cantor basins is manifest. This magnification process (i.e., zooming in on the closed left-hand interval) is

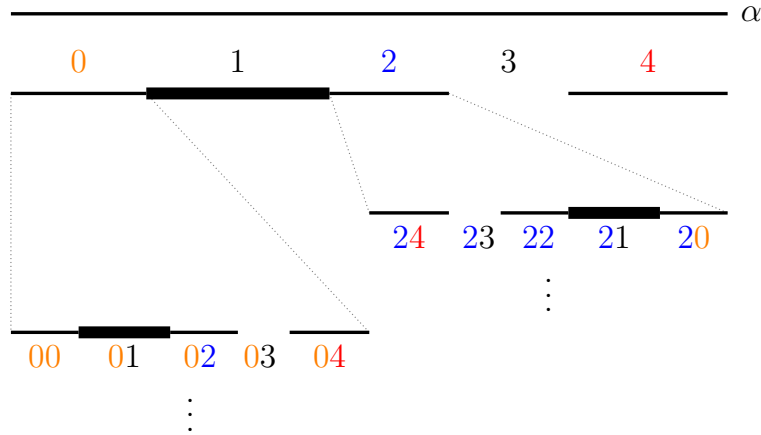


Figure 4.6. Schematic diagram which demonstrates the iterative procedure based on “decision dynamics” used to construct a one-dimensional binary black hole shadow. Beginning with a domain of the initial data α , one first removes two open intervals, corresponding to decision 1 (capture by a black hole) and decision 3 (escape to infinity). The open interval corresponding to decision 1 (depicted as a thick black line) constitutes part of the black hole shadow; the interval corresponding to decision 3 (white) does not lie in the shadow. At the next stage, one iterates on the remaining closed intervals, which comprise the initial data corresponding to trajectories whose symbolic representations start with the digit 0, 2 or 4, removing two open intervals at each stage. Iterating *ad infinitum*, we are left with an uncountable infinity of distinct points of measure zero, which correspond to the perpetual orbits.

equivalent to making the “decision” represented by the digit 0: all elements of the interval $[0, \frac{1}{3}]$ are represented in symbolic dynamics by sequences whose first digit is 0.

Constructing one-dimensional binary black hole shadows

We may use decision dynamics as a guide to develop a similar iterative process to construct the one-dimensional shadow on the initial data, defined as $\mathcal{B}_S = \mathcal{B}_+ \cup \mathcal{B}_-$ in (4.16). Denote by $\alpha_{\tilde{\mathcal{X}}}$ an initial value corresponding to a perpetual orbit, with representation \mathcal{X} in symbolic dynamics, and parity-reordered sequence $\mathcal{P}(\mathcal{X}) = \tilde{\mathcal{X}}$ (see Section 4.6.1). We need only consider the “interesting” interval $C = [\alpha_{\bar{0}}, \alpha_{\bar{4}}]$, in which all of the perpetual orbits lie. (Recall that $\alpha_{\bar{0}}$ and $\alpha_{\bar{4}}$ denote the initial values corresponding to the first and last non-escaping orbits, respectively.) From C one may remove two open intervals $O_1 = (\alpha_{0\bar{4}}, \alpha_{2\bar{0}})$ and $O_3 = (\alpha_{2\bar{4}}, \alpha_{4\bar{0}})$, which correspond to null geodesics which immediately fall into the lower black hole or escape to spatial infinity, respectively. The open interval O_1 forms part of the one-dimensional black hole shadow ($O_1 \subset \mathcal{B}_S$). Similarly, O_3 forms part of the exit basin corresponding to spatial infinity ($O_3 \subset \mathcal{B}_\infty$).

We now iterate this procedure on the remaining closed intervals. The process of iterating is equivalent to following geodesics which linger in the strong-field region until they reach the next “decision point”. This iterative procedure is demonstrated using a schematic diagram, presented in Figure 4.6.

Suppose we have iterated k times, and we are considering the closed interval $C_{\tilde{\mathcal{X}}}$, where

$\tilde{\mathcal{X}} = \mathcal{P}(\mathcal{X})$ is a parity-reordered sequence in decision dynamics of length k , which only contains *even* digits from the symbolic alphabet. Iterating involves removing the pair of open intervals $O_{\tilde{\mathcal{X}}_1} = (\alpha_{\tilde{\mathcal{X}}_1\bar{0}}, \alpha_{\tilde{\mathcal{X}}_1\bar{2}})$ and $O_{\tilde{\mathcal{X}}_3} = (\alpha_{\tilde{\mathcal{X}}_3\bar{2}}, \alpha_{\tilde{\mathcal{X}}_3\bar{0}})$. (Here, $\tilde{\mathcal{X}}\bar{0}$ represents the concatenation of the sequences $\tilde{\mathcal{X}}$ and $\bar{0}$, for example.) The subset $O_{\tilde{\mathcal{X}}_1} \subset \mathcal{B}_S$ is added to the one-dimensional shadow at each stage; the subset $O_{\tilde{\mathcal{X}}_3} \subset \mathcal{B}_\infty$ is added to the exit basin corresponding to spatial infinity. We note at this stage that $O_{\tilde{\mathcal{X}}_1}$ and $O_{\tilde{\mathcal{X}}_3}$ are the open intervals in initial data α which correspond to the set of null geodesics which go through the decision dynamics sequence \mathcal{X} , before falling into a black hole, or escaping to infinity, respectively. In the former case, one can determine the black hole into which the photon plunges as follows: count the number of digits n in the sequence \mathcal{X} that are not equal to 4; if n is even (odd), the geodesic falls into the lower (upper) black hole.

Having iterated k times, we have partitioned C into $3^k - 1$ open intervals

$$\{\{O_1, O_3\}, \{O_{01}, O_{03}, O_{21}, O_{23}, O_{41}, O_{43}\}, \dots\}, \quad (4.23)$$

which correspond to geodesics which make at most $k - 1$ decisions (represented by even digits in decision dynamics), before falling into a black hole or escaping to spatial infinity; and 3^k closed intervals

$$\{C, \{C_0, C_2, C_4\}, \{C_{00}, C_{02}, C_{04}, C_{20}, C_{22}, C_{24}, C_{40}, C_{42}, C_{44}\}, \dots\}, \quad (4.24)$$

which are made up of initial conditions corresponding to rays which linger in the strong-field region long enough to make k decisions, but whose ultimate fate has not yet been determined at this level of precision. (Here, C can be thought of as the set of all initial data corresponding to geodesics which are yet to go through a decision sequence, i.e., those which have been through an *empty* sequence.)

In the limit $k \rightarrow \infty$, one is left with an infinite number of closed sets of zero measure whose union is isomorphic to the 5-adic Cantor set; and an infinite number of open sets whose union is isomorphic to the complement of the 5-adic Cantor set in the unit interval. The binary black hole shadow is then the disjoint union of open subsets $O_{\tilde{\mathcal{X}}_1}$, i.e., the set

$$\mathcal{B}_S = \bigcup_{\mathcal{X} \in \hat{\Sigma}} O_{\tilde{\mathcal{X}}_1}, \quad (4.25)$$

where $\hat{\Sigma}$ is the set of all decision dynamics sequences \mathcal{X} of arbitrary length, which do not contain the digits 1 or 3.

At each stage of the iterative procedure outlined above, the length of the open intervals $O_{\tilde{\mathcal{X}}_1}$ and $O_{\tilde{\mathcal{X}}_3}$ as a proportion of the closed interval $C_{\tilde{\mathcal{X}}}$ from which they are removed will depend on the entire history of the geodesic; and, thus, on all preceding digits (decisions) in the trajectory's symbolic representation \mathcal{X} . In practice, the length will depend most strongly on the most recent decision taken; the dependence on previous decisions will be exponentially suppressed. We anticipate that the structure of the one-dimensional shadow will be self-similar; this will be demonstrated in Section 4.6.3.

4.6.3 Demonstrating self-similarity in the one-dimensional shadow

We now use our symbolic code as a guide to understand the fractal properties – namely self-similarity – of the one-dimensional shadow. In Figure 4.7(a), we present the one-

dimensional binary black hole shadow (i.e., exit basin diagram) for the scattering problem introduced in Section 4.4; see Figure 4.1(a) for the set-up of the problem. It is clear from Figure 4.7(a) that there are three open intervals of interest on the left (L), in the centre (C) and on the right (R). To the left of the first open interval (L), all rays escape to infinity. In between the first and second of these open intervals (L and C), all rays plunge directly into the lower black hole. All rays with initial data taken from the closed interval between the second and third open intervals (C and R) escape to infinity. Finally, all initial conditions to the right of the third open interval (R) plunge directly into the upper black hole. The highlighted intervals possess rich fractal structure. Here, we shall demonstrate that the shadow is self-similar, by magnifying each of these open intervals successively.

First, consider the leftmost open interval (L) highlighted in Figure 4.7(a). Magnifying this open interval yields Figure 4.7(b). This magnification procedure is equivalent to making the decision 0 in symbolic dynamics: all rays with initial data in this open set travel around the opposite black hole in the same sense. We see that the magnified exit basin diagram is qualitatively similar to its “parent”, shown in Figure 4.7(a). In fact, this image is identical, but the rays which plunge into the upper and lower black holes have been interchanged. We may continue in this fashion, magnifying the leftmost open interval (i.e., repeatedly making the decision 0). This yields Figures 4.7(e) and 4.7(h). Each time we magnify the left-hand interval, we obtain a new self-similar exit basin diagram but with the basins corresponding to the upper/lower black holes exchanged.

Next, we magnify the central interval (C) of the parent exit basin diagram, shown in Figure 4.7(a). This is equivalent to repeatedly making the decision 2 in symbolic dynamics. The resulting exit basin diagrams are shown in Figures 4.7(c), 4.7(f) and 4.7(i). We observe that, with each magnification, the trajectories which plunge into the upper and lower black holes are reversed (as with decision 0). Moreover, the exit basin diagram is reflected in α about the centre of the open interval, because making the decision 2 reverses the sense of the orbit. This is equivalent to reversing the parity ($P \mapsto -P$) in the parity-reordering operation \mathcal{P} (see Section 4.6.2).

Finally, we magnify the right-hand interval (R) of Figure 4.7(a), which is equivalent to making the decision 4. This yields Figures 4.7(d), 4.7(g) and 4.7(j). In this case, each magnified exit basin diagram looks the same as the parent exit basin diagram.

In Figure 4.7, we have repeatedly made the *same* decision (either 0, 2 or 4) at each stage. This process yields a particular class of self-similar exit basin diagrams, in which the proportion of the intervals corresponding to each decision are similar at each level of magnification. One could instead make *different* decisions at each stage. In this case, the resulting exit basins would again be self-similar, but the relative proportions of the intervals would depend principally on the previous digit in the decision dynamics representation of the trajectory.

We remark that the one-dimensional shadows of Figure 4.7 are qualitatively similar to the three-colour Cantor basins of Figure 4.5. In the construction of the Cantor basins, we remove one open interval from the remaining closed intervals at each step; in the construction of the one-dimensional binary black hole shadows, we remove two open intervals each

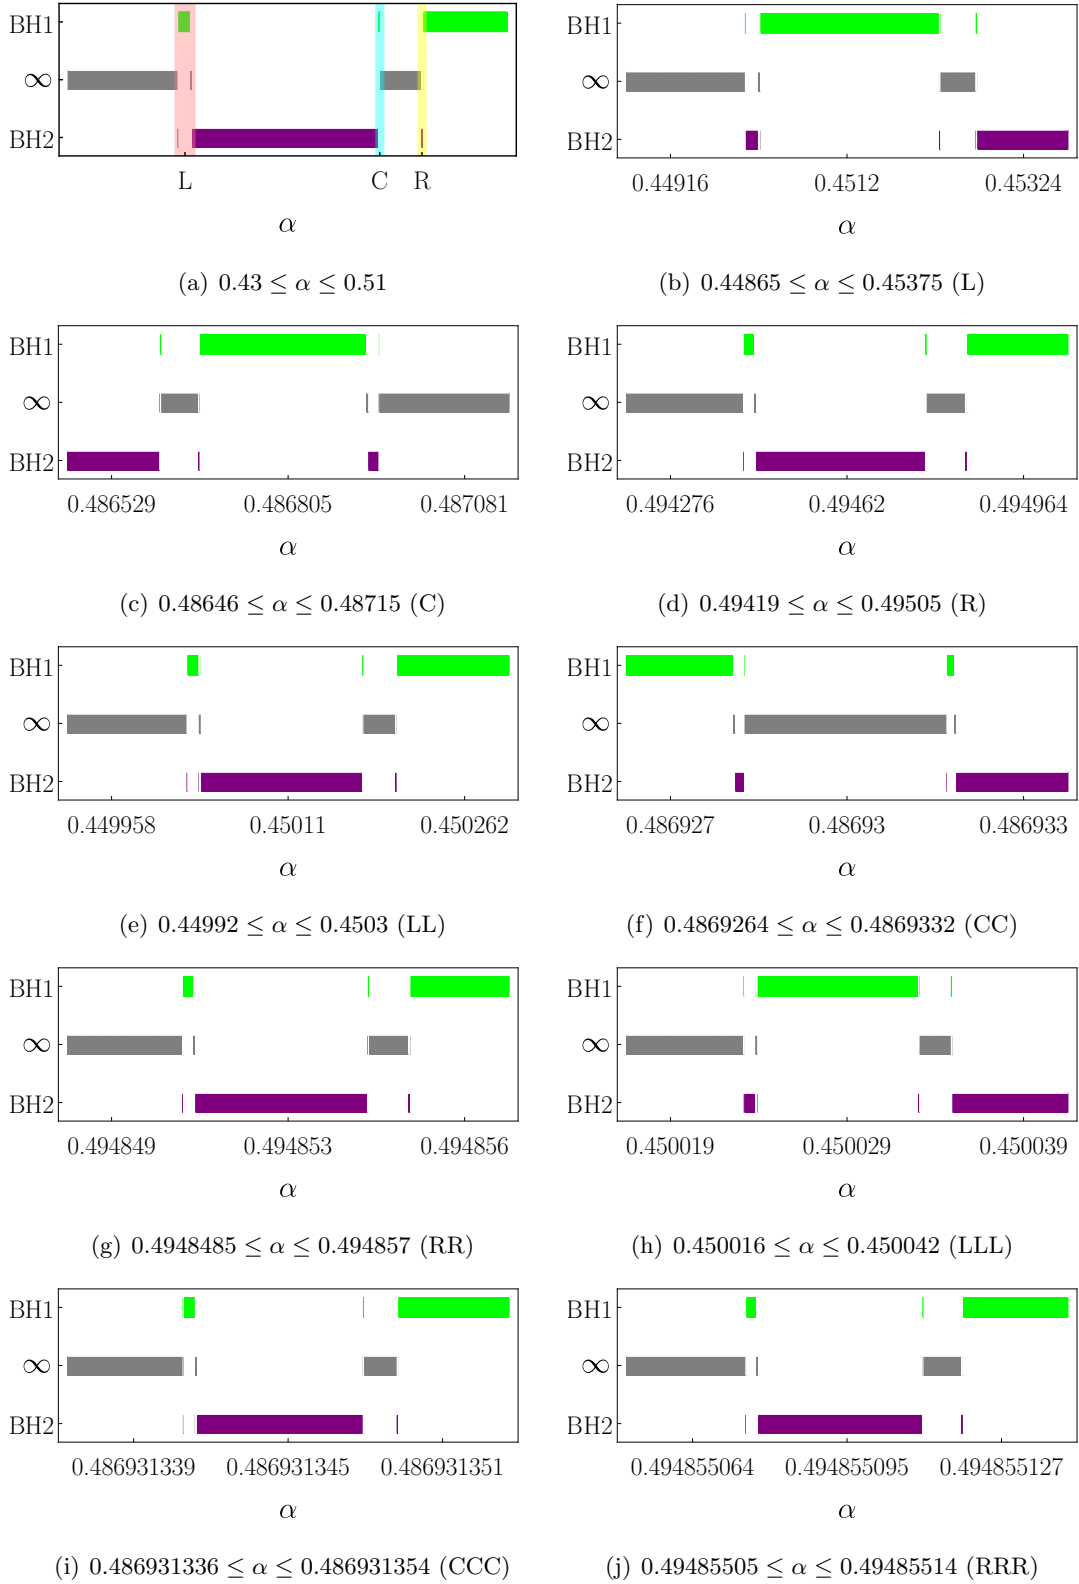


Figure 4.7. One-dimensional exit basins (shadows) for the equal-mass Majumdar–Papapetrou di-hole, with coordinate separation parameter $d = 2M$. (a) There are three regions of interest on the left (L), in the centre (C), and on the right (R), highlighted in pink, cyan, and yellow, respectively. The exit basin diagrams (b)–(j) are obtained by successively magnifying the left (L), central (C), or right (R) intervals of the “parent” exit basin diagram (a).

time we iterate. As we have seen, in both cases, the iterative procedure and the resulting self-similar basins can be understood through the use of decision dynamics.

4.6.4 The strange repellor and chaotic scattering

For a general dynamical system which evolves in time t , the set of unstable (kinematically unbounded) trajectories which remain confined to the scattering region as $t \rightarrow \infty$ constitutes the repellor Ω_R [120]; see Section 2.6.5. For example, consider the two-disc scatterer. In this model, Ω_R consists of a unique trajectory, which forever bounces between the two discs, along the straight line connecting their centres. All other trajectories scatter off the discs before reaching spatial infinity. The two-disc model exhibits *regular* dynamics, and the Kolmogorov–Sinai entropy is zero. The repellor Ω_R consists of a unique scattering singularity: the repellor is therefore *regular*.

We recall Eckhardt’s definition [119] of chaotic scattering (as stated in Section 2.6.5): scattering in a Hamiltonian system is irregular or chaotic if there exists, on some manifold of initial data, an infinity of distinct scattering singularities of measure zero which are typically arranged into a fractal set.

In contrast to the (regular) two-disc system, the repellor for the Gaspard–Rice three-disc system (see Section 4.5.2) forms a Cantor-like set: an uncountably infinite fractal set [121]. We recall that, in such cases, the repellor is called *irregular* (or *strange*). According to Eckhardt’s definition, the three-disc system therefore exhibits chaotic scattering.

In this chapter, we have demonstrated that a strange repellor Ω_R exists for planar null geodesics on the Majumdar–Papapetrou di-hole spacetime. Moreover, the arguments presented Sections 4.6.1–4.6.3 demonstrate that Ω_R is a Cantor-like set. We conclude therefore that the scattering of null rays on the Majumdar–Papapetrou di-hole spacetime is a natural exemplar of chaotic scattering: there exists a Cantor-like set of scattering singularities, corresponding to the perpetual orbits.

4.6.5 Measures of chaos

Gaspard and Rice [120] characterise the strange repellor Ω_R , and the natural measure it supports, using quantities such as Lyapunov exponents, the Kolmogorov–Sinai entropy per unit time, the Hausdorff dimension, the information dimension, the escape rate and the time-delay function. For scattering processes in general relativity, one must use such measures of chaos with caution: many of the quantities listed here are coordinate-dependent.

Here, we use the time-delay function to characterise the strange repellor Ω_R in the Majumdar–Papapetrou di-hole scattering problem. The time-delay $T(\alpha)$, shown in Figure 4.8, is defined as the coordinate time t taken by a null geodesic with initial data given by (4.15) to reach some large fixed radius $r_{\max} \gg 1$. It should be noted that $T(\alpha)$ is undefined for rays which fall into a black hole horizon, because Killing time t diverges for orbits which approach an event horizon. The time-delay $T(\alpha)$ also diverges as one approaches a scattering singularity in the initial data. This is not because it approaches a horizon, but because it asymptotes towards a perpetual orbit which lingers in the scattering region as $t \rightarrow \infty$. The time-delay function inherits the Cantor-like structure of the one-dimensional

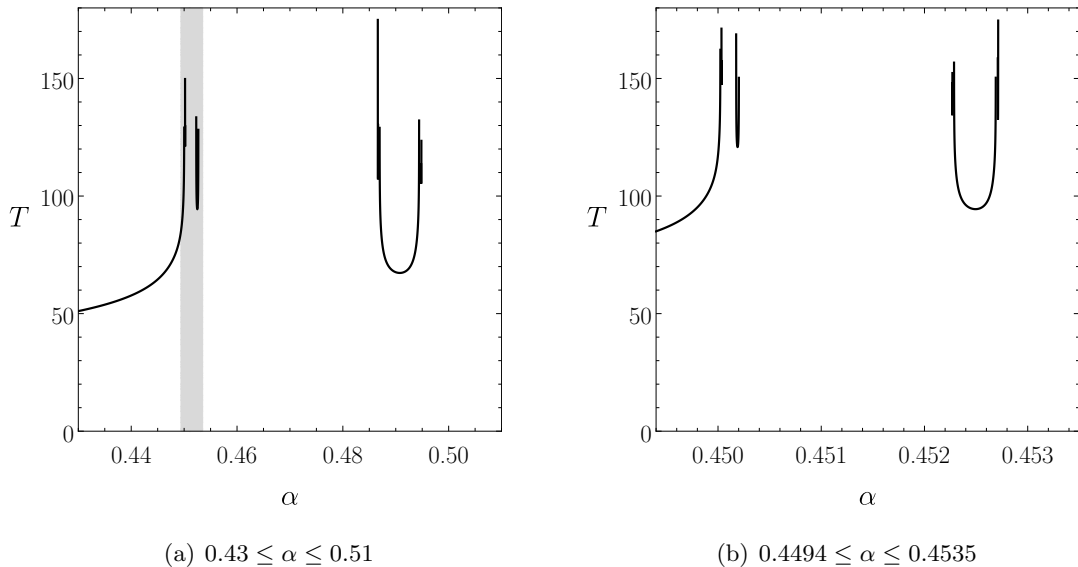


Figure 4.8. (a) Time-delay function $T(\alpha)$ for scattering by the equal-mass Majumdar–Papapetrou di-hole with initial data α . (The parameters are taken to be $M_{\pm} = M = 1$, $d = 2M$.) The time-delay is defined to be the Killing time t taken for a null ray starting at the centre of mass to escape to a large radius. The function is undefined for rays which approach a black hole horizon; $T(\alpha) \rightarrow \infty$ in the approach to a scattering singularity. The time-delay function inherits the self-similar Cantor-like structure of the strange repeller. (b) Magnification of $T(\alpha)$ which highlights the self-similarity.

shadow, due to the distribution of the scattering singularities in the initial data; this self-similarity can be seen in Figure 4.8(b), in which we magnify a domain of the time-delay function presented in Figure 4.8(a). The time-delay function is thus qualitatively similar to that used by Gaspard and Rice to characterise the strange repeller of the three-disc model (see Figure 2 of [120]).

In addition to the time-delay function, which we use here as a diagnostic for chaotic scattering, one could employ other measures to characterise the chaotic motion of (timelike and null) geodesics in the Majumdar–Papapetrou di-hole system, and related spacetimes. Here, we present a brief account of some important contributions to this area.

Dettmann *et al.* [210, 211] investigate the phase space for trajectories of charged test particles with arbitrary values of the energy and angular momentum, extending on the work of Contopoulos [140, 141]. In particular, the authors analyse exit basins for a scattering problem in the meridian plane similar to that set up in Section 4.4. They demonstrate that the boundary of the exit basins corresponding to the two black holes scales as a fractal in a diffeomorphism-invariant manner, by calculating the box-counting dimension of the boundary numerically. Moreover, the authors find that chaotic motion of *complete* timelike geodesics (i.e., those which do not plunge into a black hole) is well-described by Lyapunov exponents.

Cornish and Gibbons [213] also study geodesic motion in the field of two fixed centres (black holes) described by a family of Einstein–Maxwell–dilaton theories, which includes the Majumdar–Papapetrou di-hole (Einstein–Maxwell two-centre problem) as a special

case. The authors favour invariant methods, such as exit basins, fractal dimensions and topological entropies, over non-invariant methods, such as Lyapunov exponents and metric entropy. It is argued that geodesics will be chaotic if there exists a *chaotic invariant set* of orbits, i.e., a strange repeller (see Section 4.6.4). This set is understood as the boundary between different outcomes of the scattering problem, and the authors compute the exit basins of a scattering problem in the meridian plane, similar to this work (Section 4.4) and Dettmann *et al.* [210, 211]. Intriguingly, the results of [213] suggest that there is a transition between regular and chaotic geodesic motion as the dilaton coupling is varied. As described in Section 4.5.3, Cornish and Gibbons develop a symbolic code (“collision dynamics”) to describe the chaotic motion of null rays in the meridian plane of the two-centre Einstein–Maxwell–dilaton problem.

The numerical analysis performed by Contopoulos and others (described above) demonstrates that the motion of null geodesics on the Majumdar–Papapetrou di-hole spacetime is chaotic. Yurtsever [209] performs a geometrical analysis of the null geodesic flow on the Majumdar–Papapetrou spacetime, by reducing the problem to that of geodesic motion on a negatively curved Riemannian manifold. (Heuristically, the negative curvature is responsible for exponential divergence of trajectories in the geodesic flow – a key ingredient for chaotic motion.) The author then presents a precise formulation of chaotic motion in terms of the chaotic invariant set, before proving that the null geodesic flow of the two-centre Majumdar–Papapetrou problem satisfies this definition of chaos.

More recently, a study of the escape dynamics of photons from the field of two fixed extremal Reissner–Nordström black holes has been carried out by Alonso *et al.* [218]. Considering a scattering problem in the meridian plane (i.e., photons with $p_\phi = 0$), the authors identify the non-escaping periodic orbits which play a key role in the escape dynamics of photons from the two centres. In addition, Alonso *et al.* describe these orbits using a symbolic code similar to that of Cornish and Gibbons, and calculate characteristic quantities of chaos including the mean Lyapunov exponent of the repeller, the escape rate, the Kolmogorov–Sinai entropy, the topological entropy, the Hausdorff dimension, and the partial information dimension.

4.7 Non-planar geodesics and two-dimensional shadows

4.7.1 Non-planar rays

We now turn to non-planar motion around the two black holes, which is governed by the Hamiltonian (4.11), with a non-zero conserved azimuthal angular momentum p_ϕ . As in the case of scattering in the meridian plane (Section 4.4), it is beneficial to consider the fundamental perpetual null orbits (see Figure 4.4). Where two or more distinct but “dynamically connected” fundamental orbits exist, we anticipate chaotic scattering, and thus a Cantor-like set of scattering singularities in initial data. In this section, we shall demonstrate that this is not the only possibility when considering rays with $p_\phi \neq 0$.

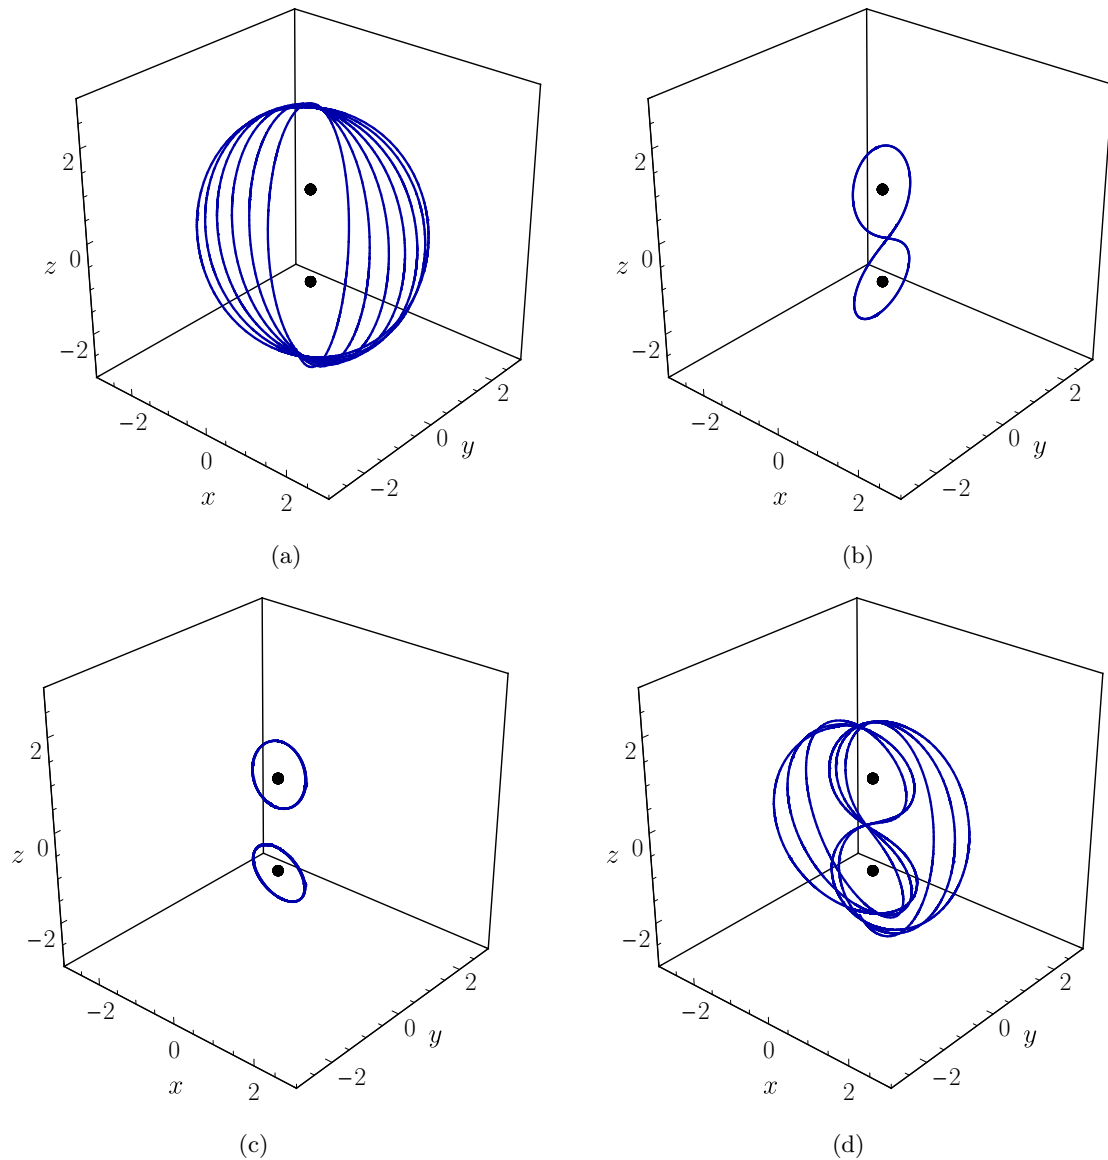


Figure 4.9. Examples of non-planar periodic null orbits with non-zero angular momentum p_ϕ , shown in \mathbb{R}^3 . We choose $M_\pm = 1$, $d = 2$ and $p_\phi = 1$ as default values. (a)–(c) Fundamental periodic orbits with symbolic representations (a) $\bar{0}$, (b) $\bar{2}$, and (c) $\bar{4}$. (d) A periodic orbit which transitions between two fundamental orbits. This orbit has symbolic representation $\bar{0}\bar{2}$ in decision dynamics. These are the non-planar versions of the fundamental orbits of Figure 4.4.




Figure 4.10. The non-planar periodic null orbits with non-zero angular momentum p_ϕ of Figure 4.9, projected onto the (ρ, z) -plane. (a)–(c) Fundamental orbits described by the decision dynamics sequences (a) $\bar{0}$, (b) $\bar{2}$, and (c) $\bar{4}$. (d) A ray which transitions between the $\bar{0}$ - and $\bar{2}$ -orbits, represented by the decision dynamics sequence $\bar{0}\bar{2}$. The grey curve is the contour $h(\rho, z) = p_\phi$. Orbits (a)–(c) touch the contour in such a way that the two-momentum is orthogonal to the contour; orbit (d) does not touch the contour.

Fundamental orbits with non-zero angular momentum

Figure 4.9 shows a selection of non-escaping orbits on the Cartesian (x, y, z) -axes in \mathbb{R}^3 for a non-zero value of the azimuthal angular momentum p_ϕ . (The parameters are chosen to be $M_\pm = M = 1$, $d = 2$ and $p_\phi = 1$.) Figures 4.9(a)–4.9(c) show the non-planar versions of the fundamental orbits of Figure 4.4, which are described by decision dynamics sequences (a) $\bar{0}$, (b) $\bar{2}$ and (c) $\bar{4}$. Figure 4.9(d) shows an example of an orbit which transitions between the fundamental orbits $\bar{0}$ and $\bar{2}$; its representation in decision dynamics is $\bar{0}\bar{2}$; see Figure 4.4(d). This case indicates that transitions between the three fundamental null orbits are still possible, despite the motion being non-planar.

In Figure 4.10, we show the orbits of Figure 4.9, projected onto the (ρ, z) -plane. Clearly, the orbits are periodic in the (ρ, z) -plane; however, it is clear from Figure 4.9 that the non-commensurate motion in ϕ means that the geodesics are not closed (in general) in three dimensions: the rays of Figures 4.9(a) and 4.9(d) trace out two-surfaces in \mathbb{R}^3 . (We only show part of the orbit for clarity.)

Analysis of the effective potential

In Section 4.3.2, we introduced the effective potential (or height function) $h(\rho, z) = \rho U^2$, which allowed us to express the geodesic potential $V(\rho, z; p_\phi)$ in a convenient factorised form. A straightforward rearrangement of the null condition $H = 0$, with Hamiltonian (4.11), yields the *energy equation*,

$$p_\rho^2 + p_z^2 = \frac{1}{\rho^2} (h + p_\phi) (h - p_\phi). \quad (4.26)$$

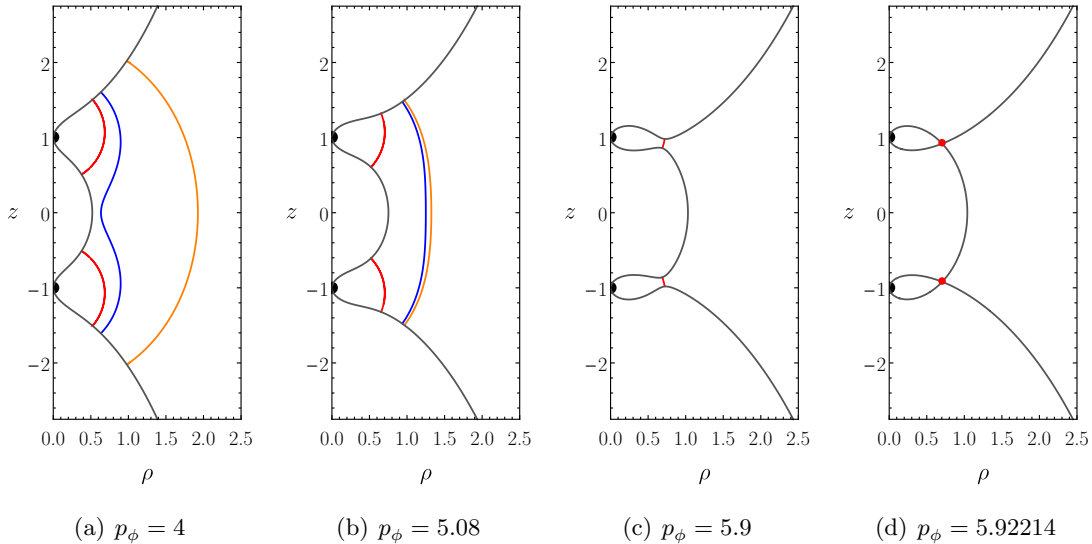


Figure 4.11. Fundamental photon orbits for the Majumdar–Papapetrou binary di-hole with coordinate separation $d = 2M$, projected onto the (ρ, z) -plane. (a) All three fundamental orbits exist, and are dynamically connected. (b) The outer orbits move closer together until they merge. (c) Only the inner pair of fundamental orbits exist. (d) The inner fundamental orbits exist up to $p_\phi = p_\phi^*$, the value of the azimuthal angular momentum for the (unstable) non-planar circular photon orbits of constant ρ and z . In each case, the contour $h(\rho, z) = p_\phi$ is shown as a grey curve.

Restricting our attention to $p_\phi \geq 0$, the effective potential h (which is clearly non-negative) determines the sign of the right-hand side of (4.26) via the factor $(h - p_\phi)$. The solutions to the equation $h(\rho, z) = p_\phi$ (i.e., the *contours* of h) are curves in the (ρ, z) -plane on which a ray may be instantaneously at rest, due to the fact that $p_\rho = 0 = p_z$ (and thus $\dot{\rho} = 0 = \dot{z}$) by the energy equation (4.26). For a given value of the azimuthal angular momentum p_ϕ , the positivity of the left-hand side of (4.26) indicates that the contour $h = p_\phi$ demarcates a *forbidden region* (corresponding to $h < p_\phi$) which the ray cannot access. The subset of the (ρ, z) -plane defined by the inequality $h \geq p_\phi$ is referred to as the *allowed region*. For a ray with $p_\phi > 0$, the forbidden region typically includes the whole symmetry axis (z -axis), with the possible exception of the black hole horizons, which are located at $z = z_\pm$.

Consider a ray with angular momentum p_ϕ . Let Γ denote the contour $h = p_\phi$ in the (ρ, z) -plane. If the ray “touches” the contour, then $\dot{\rho} = 0 = \dot{z}$ (and thus $p_\rho = 0 = p_z$), instantaneously. Hamilton’s equations for the momenta conjugate to ρ and z are

$$\dot{p}_\rho|_\Gamma = \frac{p_\phi}{\rho^2} h_{,\rho}|_\Gamma, \quad \dot{p}_z|_\Gamma = \frac{p_\phi}{\rho^2} h_{,z}|_\Gamma, \quad (4.27)$$

where we have used the fact that $h|_\Gamma = p_\phi$. As a result, rays which touch the contour Γ must have a two-momentum $(\dot{p}_\rho, \dot{p}_z)$ which is parallel to the two-vector ∇h , where $\nabla = (\partial_\rho, \partial_z)$ denotes a two-gradient operator. The vector ∇h is clearly normal to the curve $h = p_\phi$ in the (ρ, z) -plane: rays which touch the contour Γ do so in a way which is orthogonal to Γ . This can be seen in Figures 4.10(a)–4.10(c), in which the three fundamental orbits touch the contour $h = p_\phi$ orthogonally. Note that the ray in Figure 4.10(d) does not quite

touch the contour.

In Figure 4.11, we show the effect of varying p_ϕ on the contours of h and the fundamental orbits, for an equal-mass di-hole with separation parameter $d = 2$. As p_ϕ is increased, the fundamental orbits with symbolic representations $\bar{0}$ and $\bar{2}$ move closer together. When a critical value $p_\phi = \hat{p}_\phi$ is reached ($\hat{p}_\phi \approx 5.09$ for $d = 2$), the $\bar{0}$ - and $\bar{2}$ -orbits coincide (as shown in Figure 4.11(b) for $d = 2$). For $p_\phi > \hat{p}_\phi$, these two fundamental orbits cease to exist (as in Figure 4.11(c), for example). The fundamental orbits with decision dynamics representation $\bar{4}$ (light-rings around the individual black holes) persist until the contour “pinches off” at $p_\phi = p_\phi^*$ ($p_\phi^* \approx 5.92214$ for $d = 2$), as shown in Figure 4.11(d). The black holes are inaccessible to a ray incident from infinity with $p_\phi > p_\phi^*$. (Note that $\hat{p}_\phi < p_\phi^*$ for $d = 2$.)

Fixed points of the effective potential

It is clear from Figure 4.12 that the morphology of the contours of h depends on the separation between the centres d . For widely separated black holes, the Majumdar–Papapetrou di-hole system effectively behaves like a pair of isolated extremal Reissner–Nordström black holes, as shown in Figure 4.12(a). On the other hand, for a tightly bound di-hole, the system effectively resembles a single distorted black hole, as can be inferred from Figure 4.12(b). In order to better understand the system’s dependence on d , let us consider the fixed points of h and its critical contours, i.e., those which pass through the fixed points.

In Appendix B, we present a classification of the stationary points of the effective potential h both in and out of the equatorial plane. Here, we aim to deepen our understanding of the system by considering the equatorial fixed points of h , i.e., those which lie in the plane $z = 0$. (This can be achieved by considering the one-dimensional effective potential $\hat{h} = h(\rho, 0) = 1 + \frac{2}{R}$, where $R^2 = \rho^2 + \frac{d^2}{4}$.) First, define the critical values of the separation parameter

$$d_1 = \sqrt{\frac{16}{27}}M, \quad d_2 = \sqrt{\frac{32}{27}}M. \quad (4.28)$$

Using the method presented in Appendix B, we find that, for $d > d_2$, h has no stationary points in the equatorial plane. In the case $d = d_2$, there is a “cusp” in the equatorial plane at $\rho = \frac{\sqrt{5}}{2}d_2$, which corresponds to a point of inflection in the ρ -direction and a maximum in the z -direction. For separations in the range $d_1 < d < d_2$, h admits two equatorial fixed points: a saddle point at $\rho = \rho_+$, and a maximum at $\rho = \rho_-$, where $\rho_+ > \rho_- > \frac{d}{\sqrt{2}}$. For $d \leq d_1$, h has a pair of equatorial saddle points at $\rho = \rho_\pm$, where $\rho_- < \frac{d}{\sqrt{2}} < \rho_+$. In addition, we find that non-planar saddle points of h are permitted, provided that $d > d_1$.

In Figure 4.13, we show the fixed points and critical contours of the height function $h(\rho, z) = \rho U^2$ for a selection of values of the separation parameter d . The case $d = M$, shown in Figure 4.13(b), is special: h has three saddle points (one of which is in the equatorial plane), which are all connected by the *same* contour. This contour encloses a *maximum* of h . Remarkably, it is possible to find closed-form expressions for the stationary points and critical contours in terms of the golden ratio $\varphi = \frac{1}{2}(1 + \sqrt{5})$. The maximum is located at $\rho_{(1)} = \frac{\sqrt{3}}{2}M$, with $p_\phi^{(1)} = \frac{9\sqrt{3}}{2}M$. There are three saddles at $\rho_{(2)} = \frac{1}{2}5^{1/4}\varphi^{3/2}M$, $z = 0$ and $\rho_{(3)} = \frac{1}{2}5^{1/4}\varphi^{-1/2}M$, $\pm z_{(3)} = \pm \frac{M}{2\varphi}$. It is straightforward to check that all

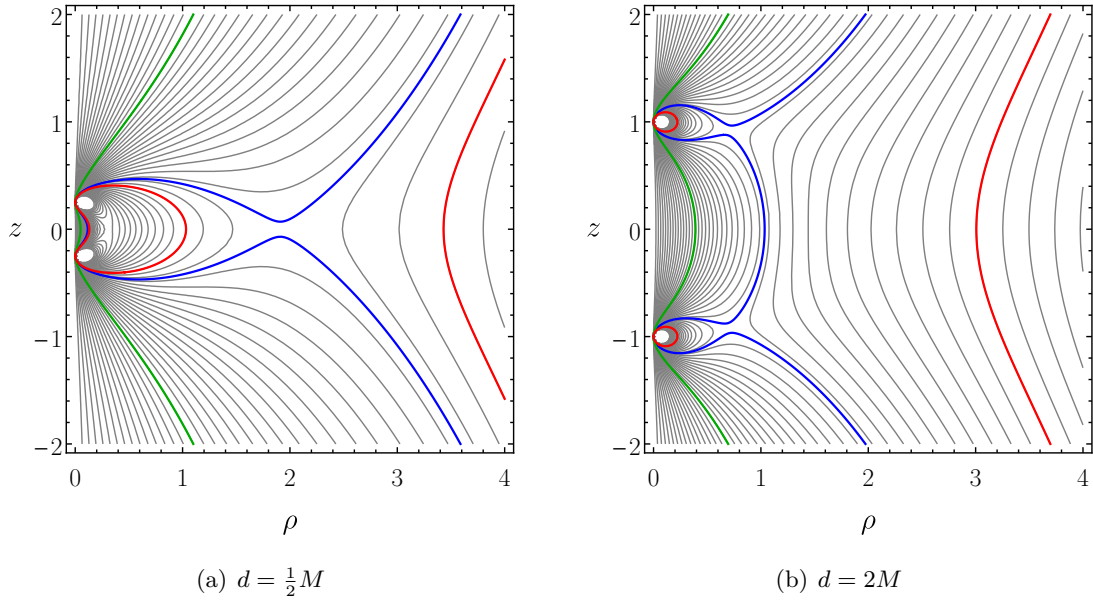


Figure 4.12. Contours of the effective potential $h(\rho, z) = \rho U^2$ for (a) $d = \frac{1}{2}M$, and (b) $d = 2M$. We highlight three typical cases: p_ϕ^A [red], p_ϕ^B [blue], p_ϕ^C [green], where $p_\phi^A > p_\phi^B > p_\phi^C$. Rays incident from infinity with $p_\phi = p_\phi^A$ are forbidden from accessing the black holes by angular momentum; rays with $p_\phi = p_\phi^B$ or $p_\phi = p_\phi^C$ are permitted to plunge into either black hole. For $p_\phi = p_\phi^B$, the ray would have to pass through a narrow channel demarcated by the contour.

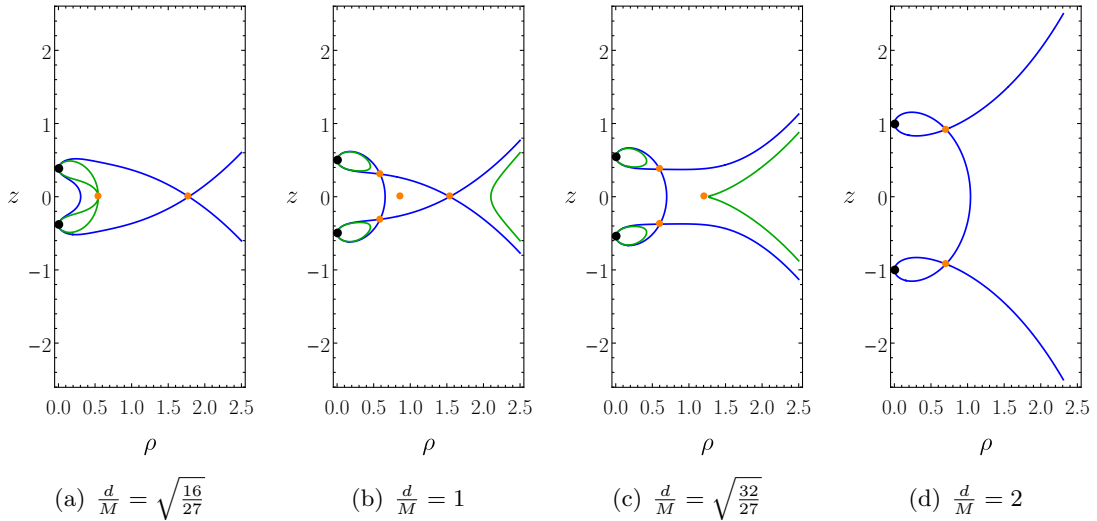


Figure 4.13. Critical contours and fixed points of the effective potential $h(\rho, z) = \rho U^2$ for a range of values of the coordinate separation d , in the equal-mass Majumdar–Papapetrou di-hole spacetime. (a) For $\frac{d}{M} \leq \sqrt{\frac{16}{27}}$, h has two saddle points in the equatorial ($z = 0$) plane. (b) For $\sqrt{\frac{16}{27}} < \frac{d}{M} < \sqrt{\frac{32}{27}}$, h has one equatorial saddle point and a pair of non-equatorial saddles; h also possesses a maximum in the equatorial plane. Stable bounded photon orbits are permitted in this regime. (c) For $\frac{d}{M} = \sqrt{\frac{32}{27}}$, there are two non-planar saddles and a “cusp” in the equatorial plane (marginal case). (d) For $\frac{d}{M} > \sqrt{\frac{32}{27}}$, h admits two saddle points out of the equatorial plane.

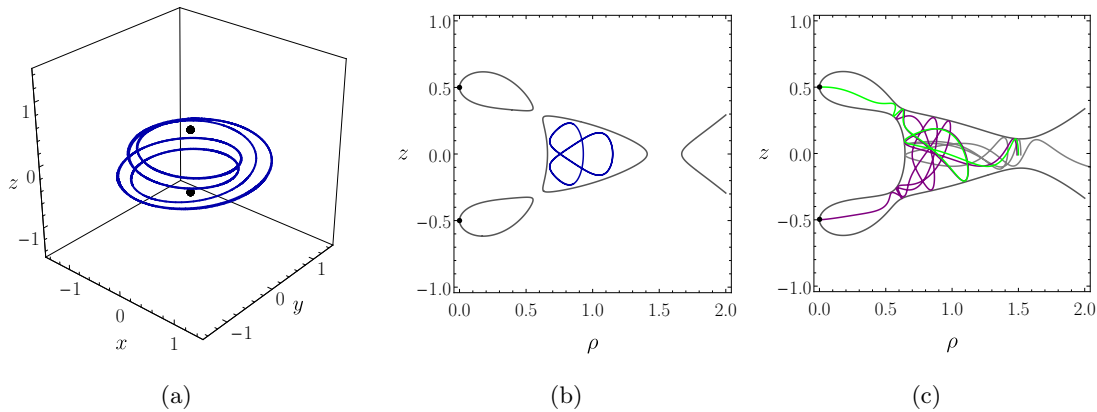


Figure 4.14. Examples of rays in the “pocket” for an equal-mass Majumdar–Papapetrou di-hole with $d = M$. (a) A bounded null geodesic with $p_\phi = 7.7 > p_\phi^*$, shown on the (x, y, z) -axes in \mathbb{R}^3 . (b) Projection of the orbit in (a) onto the (ρ, z) -plane [dark blue], with the energy contour $h(\rho, z) = p_\phi$ [dark grey]. (c) Null geodesics with $p_\phi = 7.68 < p_\phi^*$ in the “pocket” with three narrow escape channels leading to the black holes and to infinity. These initially neighbouring rays end up in different final states.

three saddle points are connected by the contour $h = p_\phi^{(2)} = p_\phi^{(3)} = \frac{1}{2}5^{5/4}\varphi^{3/2}M$. (See Appendix B for the derivation of these results.)

The saddle points of h , which correspond to unstable circular photon orbits of constant ρ and z , can be viewed as unstable “Lagrange points” for null geodesics in the Majumdar–Papapetrou di-hole system. Perturbing p_ϕ slightly from the value corresponding to a saddle point has the effect of opening a narrow channel, demarcated by neighbouring contours, in the vicinity of the saddle point. Hence, where there exist saddle points, there also exist neighbouring contours on either side of the saddle which are almost parallel to one another. Because null geodesics “touch” the contour orthogonally, we find that generic periodic unstable null orbits which “bounce” between the neighbouring parts of the contour will occur in these cases. These unstable orbits are born from perturbing the unstable “Lagrange points”. Comparing Figures 4.11(c) and 4.11(d), we see that the $\bar{4}$ fundamental orbits are born from the non-equatorial saddle points in the case $d = 2M$.

Bounded null geodesics

The existence of a maximum of the effective potential h in the regime $d_1 < d < d_2$ implies the existence of *bounded* null geodesics. These are confined to a compact subregion of the (ρ, z) -plane. Using the relationship between the geodesic potential $V(\rho, z)$ and the effective potential $h(\rho, z)$ given in (4.11), one can show that a local maximum of h corresponds to a local minimum of the potential V ; hence, the circular orbit of constant ρ and z corresponding to the maximum of h is *stable*. The existence and phenomenology of stable null geodesics will be explored in the context of stationary axisymmetric electrovacuum spacetimes in Chapter 6.

Figure 4.14 shows an example of a bounded null geodesic for a di-hole with $d = M$. In Figure 4.14(a), we show the orbit in \mathbb{R}^3 on a set of Cartesian axes; the orbit is clearly

confined to a toroidal region around the two black holes. Figure 4.14(b) shows the orbit’s projection onto the (ρ, z) -plane and the corresponding contour $h = p_\phi$; the periodicity of the orbit is manifest.

Decreasing p_ϕ from p_ϕ^* , it is possible to construct a “pocket” in the (ρ, z) -plane which connects the scattering region to the two black holes and to spatial infinity via three narrow channels; see Figure 4.14(c). The pocket and the escape channels are demarcated by the contour $h = p_\phi^* - \Delta p_\phi$, where $\Delta p_\phi \gtrsim 0$ is a small perturbation. As $\Delta p_\phi \rightarrow 0$, the width of the escapes goes to zero. In Figure 4.14(c), we show three initially neighbouring rays which begin in the “pocket” (inside the scattering region), but which end up in different final states.

4.7.2 Binary black hole shadows

Ray-tracing set-up and numerical method

In the section, we revisit the definition of a black hole shadow (see Section 3.3.4). Bohn *et al.* [190] employ a ray-tracing approach, and define a shadow as “a region of the [observer’s] image where geodesics are traced backwards in time from the camera to the black hole”. The camera provides a natural one-to-one correspondence between a “pixel” on a two-dimensional black hole shadow image and a null geodesic.

Our primary aim is to study the structure of binary black hole shadows for the Majumdar–Papapetrou di-hole geometry. To realise the black hole shadows, we trace rays which pass orthogonally through a planar surface (the image plane) with centre $(x, y, z) = (x_0, 0, z_0)$ in Cartesian coordinates, where $r_0 = \sqrt{x_0^2 + z_0^2}$ is the distance from the centre of the surface to the origin (taken to be the centre of mass). In our computations, we typically take $r_0 = 50M$. The angle of incidence θ is defined via the equations $\sin \theta = \frac{x_0}{r_0}$, $\cos \theta = -\frac{z_0}{r_0}$; that is, θ is measured with respect to the negative z -axis in the anti-clockwise sense.

The observer’s two-dimensional image plane is labelled by a local system of Cartesian-type coordinates (X, Y) . We determine whether the point (X, Y) is in the shadow by integrating Hamilton’s equations backwards in time, taking as initial data

$$x(0) = x_0 + X \cos \theta, \quad y(0) = Y, \quad z(0) = z_0 - X \sin \theta, \quad (4.29)$$

$$p_x(0) = -U_0^2 \sin \theta, \quad p_y(0) = 0, \quad p_z(0) = U_0^2 \cos \theta, \quad (4.30)$$

where $U_0 = U(x, y, z)|_{\lambda=0}$. If the ray approaches a black hole horizon, then the point (X, Y) is assumed to be in the shadow. If the ray approaches spatial infinity, then the point (X, Y) is in the non-shadow region.

As described in Section 4.3.3, geodesic motion on the Majumdar–Papapetrou di-hole spacetime is non-integrable. We therefore resort to numerical methods to evolve the equations of motion. In particular, we perform ray-tracing by numerically evolving Hamilton’s equations subject to the initial data (4.29)–(4.30). This is achieved by employing Mathematica’s `NDSolve` function.

In practice, the shadow is realised using a finite-resolution grid of pixels on the image plane. A *pixel* is defined to be a square region of the image plane with side length

L and centre (X, Y) . The pixel with midpoint (X, Y) is in the black hole shadow if and only if the null geodesic with initial data given by (4.29)–(4.30) approaches a black hole horizon. Practically, this involves implementing one of two halting conditions. Our numerical integration is stopped when either

$$x^2(\lambda) + y^2(\lambda) + (z(\lambda) - z_{\pm})^2 \leq \varepsilon, \quad 0 < \varepsilon \ll 1; \quad (4.31)$$

or when

$$x^2(\lambda) + y^2(\lambda) + z^2(\lambda) > r_0. \quad (4.32)$$

The former case indicates absorption by the black hole with horizon at $z = z_{\pm}$; the latter case indicates scattering.

In the approach outline above, the image plane defines a set of local observers at each point (X, Y) . With a change of emphasis, one can instead define a black hole shadow with respect to a single observer, by performing ray-tracing from a single point in spacetime, varying the initial elevation and azimuth of the ray’s two-momentum. The two definitions are essentially equivalent as $r_0 \rightarrow \infty$.

Gallery of two-dimensional shadows

In Figure 4.15, we present a gallery of two-dimensional binary black hole shadows cast by the equal-mass ($M_{\pm} = M = 1$) Majumdar–Papapetrou di-hole system with coordinate separation $d = 2$, for a selection of values of the angle of incidence θ . The shadows are realised by using the ray-tracing approach described above. Pixels which belong to the shadow of the upper (lower) black hole are coloured green (purple); those which correspond to spatial infinity are coloured white. We observe that, as anticipated from the simulations by Bohn *et al.* [190], the binary shadow is *not* simply the superposition of two singleton black hole shadows. Rather, each black hole has a primary shadow – either topologically equivalent to a disc or an annulus – surrounded by a hierarchy of disconnected subsidiary components.

Let us consider the effect of increasing θ in Figure 4.15. For $0^\circ \leq \theta \leq 40^\circ$, the shadow is distorted but remains qualitatively similar to the $\theta = 0^\circ$ case. When $\theta \sim 50^\circ$, we clearly see eyebrow-like features [190] for the first time: the secondary annular shadow corresponding to the lower black hole (depicted in purple) has fragmented into two disconnected arc-like components, as shown in Figure 4.15(f). When we reach $\theta \sim 70^\circ$, the primary annular shadow corresponding to the upper black hole (plotted in green) splits into two disconnected components – one of which forms the primary globular shadow; the other forms a secondary eyebrow – as shown in Figure 4.15(h).

When $\theta = 90^\circ$, the observer has a “side-on” view of the black holes. In this case, we observe two primary shadows of equal size and shape. The primary shadow is surrounded by a hierarchy of self-similar eyebrows, which have a Cantor-like structure. This case was shown in Figure 3(b) of [217] and Figure 2 of [199].

One could form a “flip book” animation using Figures 4.15(a)–4.15(j). This would mimic the effect of an observer passing by a pair of fixed black holes; or a static observer viewing a pair of rigidly rotating black holes.

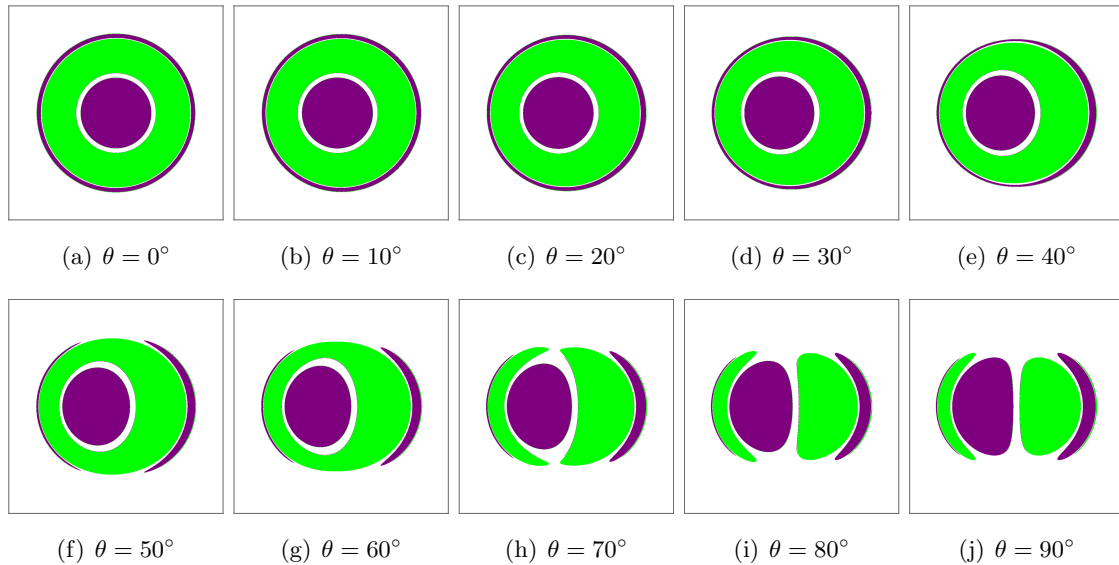


Figure 4.15. Gallery of two-dimensional binary black hole shadows for the equal-mass ($M_{\pm} = M$) Majumdar–Papapetrou di-hole with coordinate separation $d = 2M$. Initial conditions corresponding to rays which plunge into the upper (lower) black hole are presented in green (purple).

On-axis case

Let us consider the $\theta = 0^\circ$ case, shown in Figure 4.15(a), in more detail. This depicts the Majumdar–Papapetrou di-hole shadow as seen by an observer located on the symmetry axis. Figure 4.16(a) shows a one-parameter scattering problem, in which rays are fired towards the black holes from infinity, initially parallel to the symmetry axis, with impact parameter b . In Figure 4.16(b), we show the two-dimensional shadow in the image plane. By symmetry, this shadow may be constructed from the “area of revolution” of the one-dimensional shadow of Figure 4.16(c), which depicts the fate of the rays as a function of the *impact parameter* b , which is the perpendicular distance from the initial position of the ray to the symmetry axis.

Let us analyse the crude features of the shadow as b is increased from zero. For $b \sim 0$, all rays plunge directly into the lower black hole. Such rays are responsible for the primary disc-shaped shadow of the lower black hole, depicted in purple. Then, there is an interval in initial data corresponding to rays which pass between the two black holes and escape to infinity. These rays correspond to the “gap” between the primary shadows. Then, there is an open interval in b corresponding to rays which plunge directly into the upper black hole. This open interval in initial data corresponds to the primary annular shadow of the upper black hole, shown in green. Finally, for sufficiently large values of b , all rays escape to infinity.

This crude analysis overlooks the self-similar features of the on-axis shadow. A more careful analysis may be performed with the aid of symbolic dynamics, introduced in Section 4.5.3. The outer edge of the primary shadow of the lower black hole (i.e., the purple disc in the centre of the image) corresponds to the fundamental orbit with decision dynamics representation $\bar{4}$. Similarly, the outer edge of the shadow corresponds to the fundamental

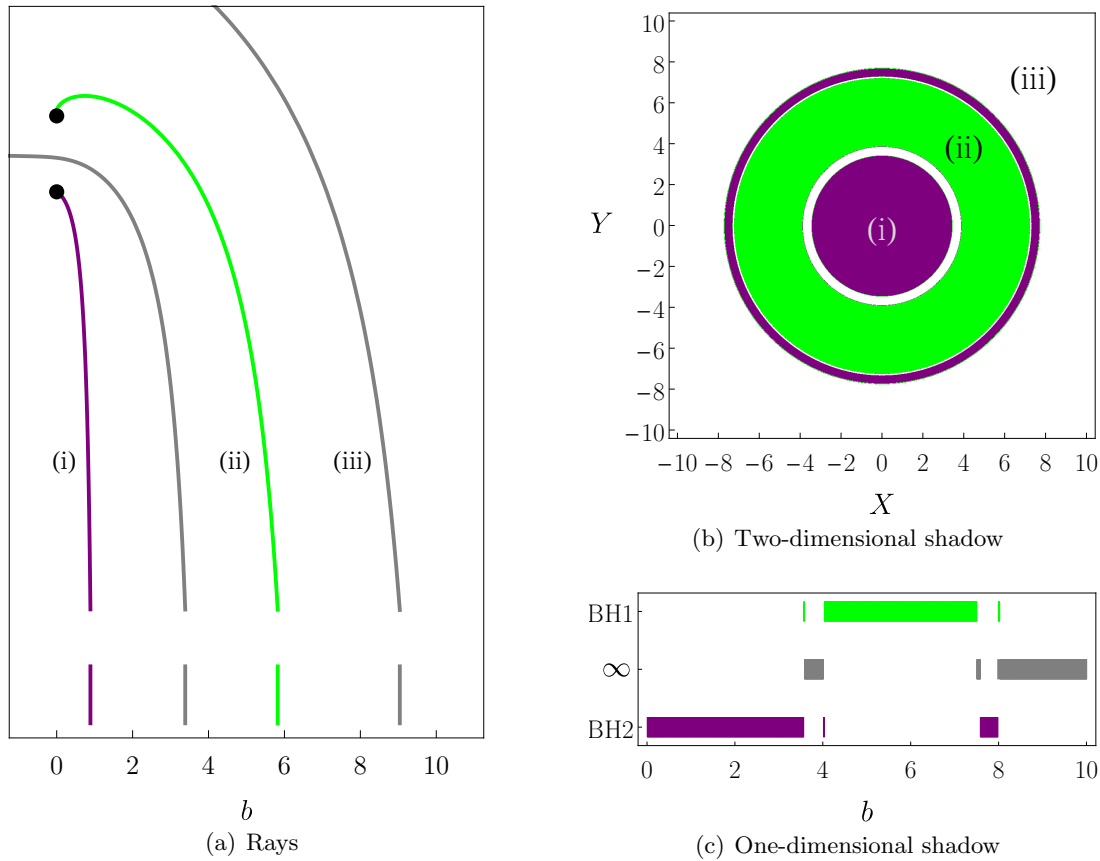


Figure 4.16. On-axis shadow and the impact parameter. (a) Null rays on an initial data surface labelled by the impact parameter b . (b) Two-dimensional shadow viewed by an observer on the symmetry axis. (c) One-dimensional shadow for rays with impact parameter b . Rays can (i) fall directly into the lower black hole [purple]; (ii) plunge directly into the upper black hole [green]; (iii) escape to infinity [grey/white]. In between these possibilities, we observe rich dynamics which results in self-similar shadows.

orbit which is represented in decision dynamics by the sequence $\bar{0}$. As described in Section 4.6, there exists an uncountable infinity of impact parameters $b_{\mathcal{X}}$ which correspond to the non-escaping perpetual orbits. Here, \mathcal{X} is a sequence in decision dynamics of infinite length which does not contain the digits 1 or 3. The ordering of $b_{\mathcal{X}}$ in initial data is determined by the ordering function F , defined in Section 4.6.1; however, the ordering is now reversed, so we begin with parity $P = -1$ in the parity-reordering operation \mathcal{P} . The one-dimensional shadow on initial data b may be constructed iteratively using symbolic dynamics, as in Section 4.6.2. Hence, the one-dimensional shadow has a Cantor-like structure which is inherited by the two-dimensional shadow due to the symmetry of the on-axis scattering problem.

Side-on case

We now turn our attention to the case $\theta = 90^\circ$, in which an observer views the system from “side-on”, and the centre of the image plane is in the equatorial ($z = 0$) plane. In this

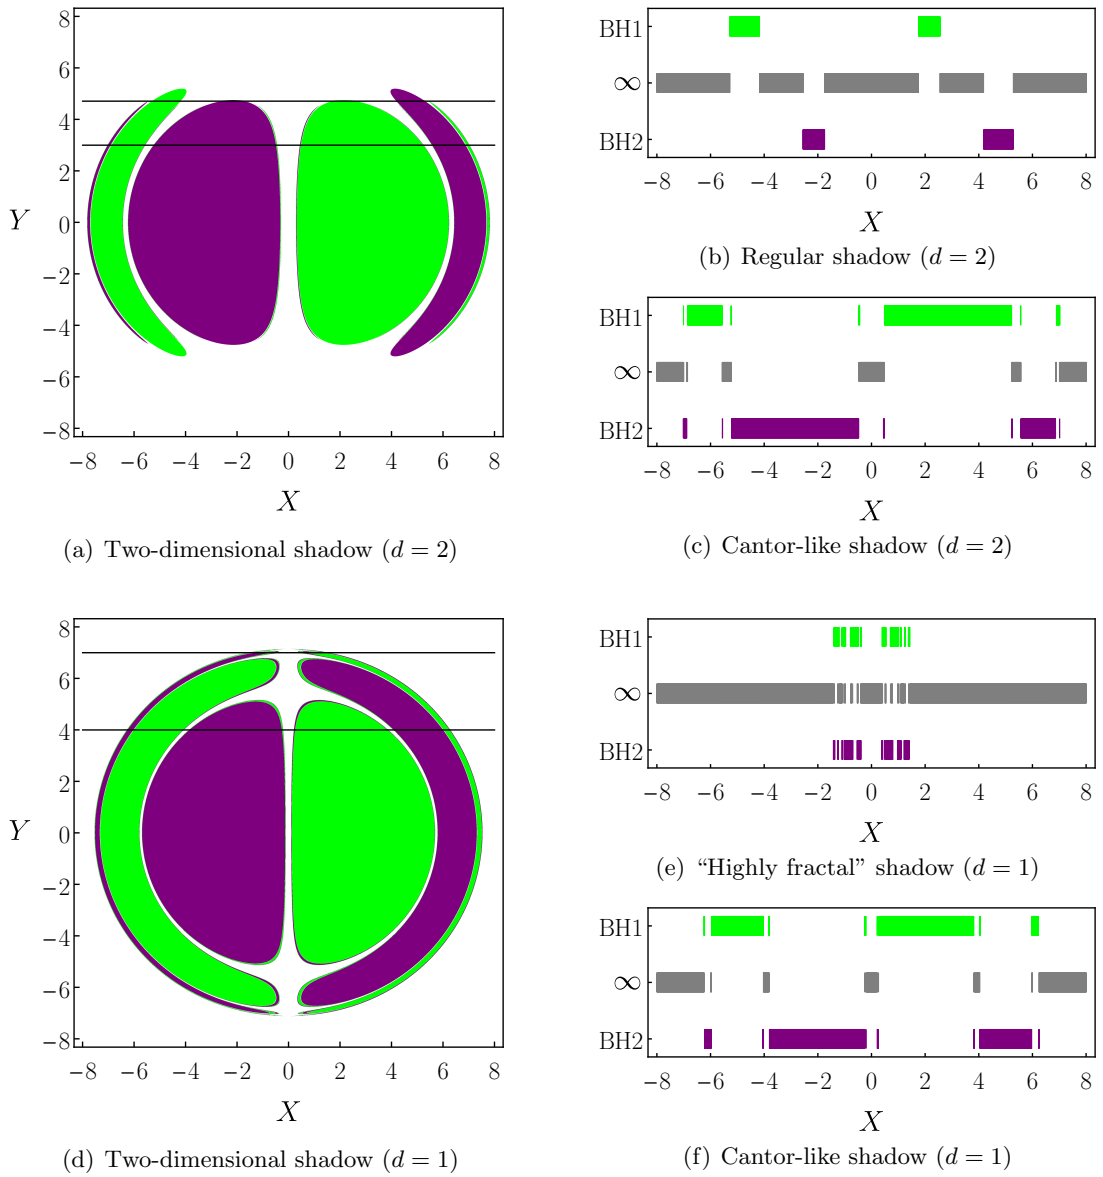


Figure 4.17. Fractal structure of the two-dimensional shadows of the equal-mass ($M_{\pm} = M = 1$) Majumdar-Papapetrou di-hole, shown in the (X, Y) -plane. The two-dimensional shadow can be viewed as the union of one-dimensional shadows for a fixed value of p_{ϕ} (which corresponds to a fixed value of Y). (a) Two-dimensional shadow image for $d = 2$. (b) One-dimensional slice with $d = 2$ at $Y = 4.7$, which exhibits regular structure due to the non-existence of the $\bar{0}$ and $\bar{2}$ fundamental orbits. (c) Cantor-like one-dimensional slice at $Y = 3$ for $d = 2$. (d) Two-dimensional shadow for $d = 1$. (e) One-dimensional shadow for $d = 1$ with $Y = 7$ which exhibits "highly fractalised" structure, due to the existence of a scattering region with three narrow "throats"; see Figure 4.14(c). (f) Cantor-like one-dimensional slice at $Y = 4$ for $d = 1$.

case, the two-dimensional black hole shadow image, with image plane coordinates (X, Y) , can be viewed as a set of one-dimensional shadows, each given by a fixed value of the conserved azimuthal angular momentum p_ϕ . To see this, consider the change of spacetime coordinates $x^a \mapsto x^{a'}$. Under such a transformation, the momenta transform according to $p_a \mapsto p_{a'} = \frac{\partial x^a}{\partial x^{a'}} p_a$. The standard relationship between Cartesian coordinates $\{x, y, z\}$ and cylindrical polar coordinates $\{\rho, z, \phi\}$ then yields the relationship $p_\phi = xp_y - yp_x$. A photon with initial momentum normal to the image plane has $p_y(0) = 0 = p_z(0)$ and $p_x(0) = -U_0^2$, by the Hamiltonian constraint. Hence, the azimuthal angular momentum can be related to the image plane coordinates via $p_\phi = YU_0^2$. Moreover, for an observer in the far-field ($r_0 \rightarrow \infty$), we have $U_0 \rightarrow 1$. Hence, a scattering problem with $p_\phi = \text{constant}$ admits a one-dimensional shadow, which corresponds to a curve $Y = \text{constant}$ in the two-dimensional shadow image. This is shown in Figure 4.17.

Consider first the case $d = 2$, shown in Figures 4.17(a)–4.17(c). One-dimensional slices corresponding to sufficiently small values of p_ϕ exhibit self-similar Cantor-like structure, as shown in Figure 4.17(c). However, for large values of p_ϕ , the one-dimensional slices (e.g. close to the top of the shadow) do *not* exhibit fractal structure. This qualitative difference in shadow structure was anticipated in the analysis of the non-planar fundamental photon orbits in Section 4.7.1. In Figure 4.11(c), we show a value of the angular momentum p_ϕ for which the fundamental orbits $\bar{0}$ and $\bar{2}$ are forbidden, but for which the pair of fundamental orbits with representation $\bar{4}$ exist, allowing absorption by the black holes. In this regime, dynamical transitions between fundamental orbits are not possible, and chaotic scattering does not occur; however, the shadow is non-empty because absorption is permitted. As a result, the one-dimensional shadows will be regular (not Cantor-like), as illustrated in Figure 4.17(b).

In the case $d = 1$, the one-dimensional scattering problem for values of p_ϕ close to p_ϕ^* is associated with the existence of a “pocket” feature with three narrow escape channels; see Figure 4.14(c). The pocket acts as a “randomising” region, and the orbits which enter are highly chaotic. The resulting one-dimensional shadows are therefore “highly fractalised”, as can be seen in Figure 4.17(e).

The analysis carried out in this section demonstrates that there is clearly a qualitative difference between the shadows for $d = 1$ and $d = 2$. This could be anticipated from the analysis of the effective potential h ; see Section 4.7.1, in particular Figures 4.12 and 4.13. In particular, we observed a “phase change” in the behaviour of null geodesics as the separation was varied. This will be explored in more depth using quantitative methods in Chapter 5.

4.8 Following rays through the event horizons

In this section, we consider the consequences of following null geodesics *through* the event horizons in the maximally extended Majumdar–Papapetrou di-hole spacetime [204]. As we shall show, following geodesics in this way results in richer chaotic phenomena.

We recall that the “points” $(x, y, z) = (0, 0, z_\pm)$ are *not* curvature singularities but coordinate singularities at which $U \rightarrow \infty$ [204]. At these points, the coordinate time t

diverges ($t \rightarrow \infty$); however, the affine parameter for each geodesic remains finite, as do curvature invariants such as the Kretschmann scalar $R_{abcd}R^{abcd}$. (The “points” at which $U \rightarrow \infty$ are actually null *surfaces* of finite area with topology $\mathbb{R} \times S^2$, which correspond to the *event horizons* of the two black holes.)

4.8.1 Extremal Reissner–Nordström case

Before considering the behaviour of null geodesics which pass through the black hole event horizons in the Majumdar–Papapetrou di-hole spacetime, it will be beneficial to review the singleton black hole case, i.e., the (maximally extended) extremal Reissner–Nordström spacetime. The spacetime metric in isotropic spherical polar coordinates $\{t, r, \theta, \phi\}$ takes the form

$$ds^2 = -\frac{dt^2}{U^2} + U^2 (dr^2 + r^2 d\Omega^2), \quad U(r) = 1 + \frac{M}{r}, \quad (4.33)$$

where $d\Omega^2 = d\theta^2 + \sin^2 \theta d\phi^2$ is the line element on the unit two-sphere. The point $r = 0$ in isotropic coordinates corresponds to a null surface of finite area with topology $\mathbb{R} \times S^2$ – the event horizon of the black hole.

The metric (4.33) may be transformed into a more familiar form via the change of coordinates $\hat{r} = r + M$, so that the line element in standard (Schwarzschild-type) coordinates $\{t, \hat{r}, \theta, \phi\}$ reads

$$ds^2 = -f(\hat{r}) dt^2 + \frac{dr^2}{f(\hat{r})} + \hat{r}^2 d\Omega^2, \quad f(\hat{r}) = \left(1 - \frac{M}{\hat{r}}\right)^2. \quad (4.34)$$

One may instead choose to replace the temporal coordinate t with an ingoing (–) or outgoing (+) null coordinate $w_{\pm} = t \pm F(r)$, where $F'(r) = U(r)$. Performing this coordinate transformation leads to the line element

$$ds^2 = g_{ab} dx^a dx^b = -\frac{dw_{\pm}^2}{U^2} \pm dw_{\pm} dr + U^2 r^2 d\Omega^2. \quad (4.35)$$

These coordinates are similar to the familiar Eddington–Finkelstein coordinates [61, 228] for the Schwarzschild black hole. In the $\{w_{\pm}, r, \theta, \phi\}$ coordinate system, the metric components g^{ab} are *regular* at the horizon $r = 0$. One may therefore employ the Hamiltonian $H = \frac{1}{2} g^{ab} p_a p_b$ with the ingoing (outgoing) null coordinate w_- (w_+) to follow null rays into (out of) the black hole horizon.

In Figure 4.18, we present a Penrose–Carter diagram [229, 230] which illustrates the causal structure of a (singleton) extremal Reissner–Nordström black hole of mass M . In Schwarzschild-type coordinates (4.34), the black hole horizon is located at $\hat{r} = M$, and the curvature singularity is located at $\hat{r} = 0$. In isotropic coordinates (4.33), the event horizon is located at $r = 0$. Null geodesics may be evolved through the black hole horizon from Region I ($r > 0$) to Region III ($r < 0$) by switching to *ingoing* Eddington–Finkelstein-type coordinates. All rays inside Region III which have non-zero angular momentum will avoid the timelike singularity, and hence will emerge through the white hole horizon into a new asymptotically flat spacetime (Region I'). An example of a ray which has been extended through the black hole horizon, and then through the white hole horizon into a new asymptotically flat spacetime is shown in Figure 4.18.

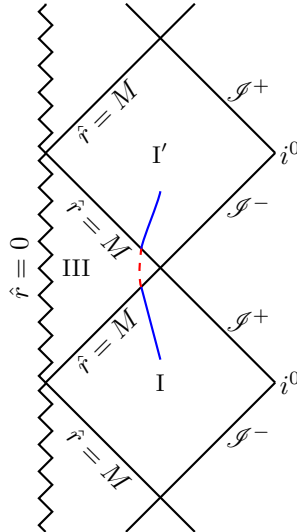


Figure 4.18. Penrose–Carter diagram for the extremal Reissner–Nordström black hole. The diagram includes an example of a null geodesic which is emitted by a source in Region I, then crosses the black hole horizon from Region I to Region III, before emerging from a white hole horizon in a new asymptotically flat spacetime, Region I'. The part of the trajectory which is evolved using Eddington–Finkelstein-type coordinates is shown as a red dashed curve. The transition from ingoing to outgoing null coordinates occurs at the turning point, where $\frac{d\hat{r}}{d\lambda} = 0$.

4.8.2 Majumdar–Papapetrou di-hole case

In the case of the Majumdar–Papapetrou di-hole, one may track ingoing/outgoing null rays as follows. Suppose that we wish to follow a null geodesic through the event horizon of the lower black hole. We transform from standard isotropic coordinates to a spherical coordinate system $\{t, r, \theta, \phi\}$ centred on the event horizon of the lower black hole $(x, y, z) = (0, 0, z_-)$. In this coordinate system, the line element is given by

$$ds^2 = -\frac{dt^2}{U^2} + U^2 (dr^2 + r^2 d\Omega^2), \quad U(r, \theta) = 1 + \frac{M_-}{r} + \frac{M_+}{\sqrt{r^2 - 2dr \cos \theta + d^2}}. \quad (4.36)$$

We introduce the radial function $V(r) = 1 + \frac{M_-}{r} + \frac{M_+}{d}$ and the outgoing/ingoing null coordinate $w_{\pm} = t \pm F(r)$, such that $F'(r) = V(r)$. In coordinates $\{w_{\pm}, r, \theta, \phi\}$ the line element reads

$$ds^2 = g_{ab} dx^a dx^b = -\frac{dw_{\pm}^2}{U^2} \pm \frac{2V^2}{U^2} dw_{\pm} dr + \frac{U^4 - V^4}{U^2} dr^2 + U^2 r^2 d\Omega^2, \quad (4.37)$$

and the corresponding Hamiltonian is given by

$$H = \frac{1}{2U^2} \left[-(U^4 - V^4) p_{w_{\pm}}^2 \pm 2V^2 p_{w_{\pm}} p_r + p_r^2 + \frac{1}{r^2} p_{\theta}^2 + \frac{1}{r^2 \sin^2 \theta} p_{\phi}^2 \right]. \quad (4.38)$$

The momenta are related by $p_{w_{\pm}} = p_t \pm V(r)p_r$. We note here that

$$\lim_{r \rightarrow 0} \frac{U^4 - V^4}{U^2} = \frac{4M_+ M_- \cos \theta}{d^2}, \quad \lim_{r \rightarrow 0} \frac{V^2}{U^2} = 1, \quad \lim_{r \rightarrow 0} \frac{1}{U^2 r^2} = \frac{1}{M_-}, \quad (4.39)$$

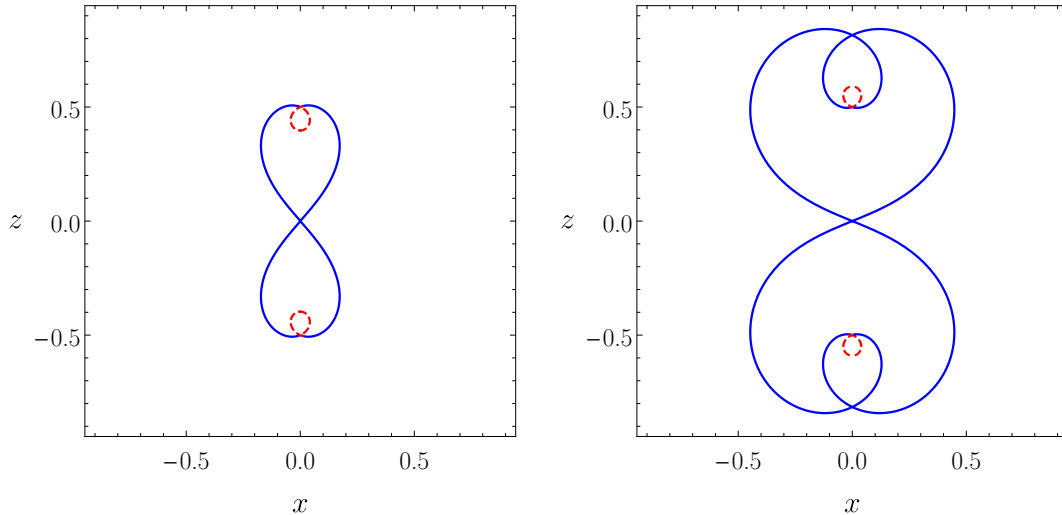


Figure 4.19. Examples of periodic null geodesics which pass through the event horizons of both black holes in the Majumdar–Papapetrou di-hole spacetime. The sources are separated by coordinate distance $d = 1$. Rays which emanate from the centre of mass between the black holes fall into the upper black hole. The geodesics are then tracked using the ingoing ($-$) and outgoing ($+$) null coordinates w_{\pm} ; the switch from ingoing to outgoing coordinates occurs at the symmetry point (i.e., where the ray crosses the z -axis).

so the Hamiltonian (4.38) is non-singular as we approach the event horizon of the lower black hole, $r = 0$. We may therefore evolve Hamilton’s equations through the coordinate singularity at $r = 0$, using the outgoing/ingoing null coordinate w_{\pm} .

In practice, we may follow a ray which begins in the exterior asymptotically flat region of the Majumdar–Papapetrou di-hole spacetime, until it reaches the vicinity of a black hole horizon. The null geodesic may then be evolved through the horizon (from $r > 0$ to $r < 0$) using a coordinate patch which is centred on the relevant black hole and a null ingoing coordinate w_- . All rays with non-zero angular momentum will avoid the timelike singularity; they will therefore reach a turning point at which $\dot{r} = 0$. At this point, one may switch from the ingoing to the outgoing null coordinate $w_- \mapsto w_+$, switching also the radial momenta $p_r^- \mapsto p_r^+ = p_r^- - 2V^2 p_t$. One may track the null ray from the interior region to a new asymptotically flat exterior region (through a white hole horizon) from $r < 0$ to $r > 0$, before switching back to isotropic coordinates. This process may be repeated as necessary, for all rays which plunge into a black hole horizon and emerge through a white hole horizon.

Figure 4.19 shows examples of *periodic* null geodesics which pass through the black holes’ event horizons in the Majumdar–Papapetrou di-hole spacetime. (We choose $d = 1$ as the coordinate separation between the holes.) Parts of the trajectory which are evolved in isotropic coordinates are shown in blue; the part of the geodesic which was evolved using the null coordinates w_{\pm} is depicted as a red dashed line. The switch from ingoing ($-$) to outgoing ($+$) coordinates occurs at the symmetry point. The exterior ($r > 0$) and interior ($r < 0$) are presented on the same plot by only plotting the magnitude of r in each case, i.e., we plot the coordinates $(x, z) = (|r| \cos \theta, |r| \sin \theta)$. The figure is slightly

misleading: each time the photon passes through a white hole horizon, it emerges into a *new* asymptotically flat exterior region. All of these regions are overlaid in Figure 4.19.

4.9 Discussion

4.9.1 Extensions

Chaotic lensing in other spacetimes

At the beginning of this chapter, the Majumdar–Papapetrou di-hole spacetime was introduced as a surrogate model for a dynamical binary black hole system in the late stages of its inspiral, before merger. We caution that the Majumdar–Papapetrou di-hole fails to be a “physically realistic” surrogate in a number of ways. Firstly, the components of a realistic binary (and indeed all astrophysical black holes) are thought to have negligible charge. (Recall that each black hole in the Majumdar–Papapetrou system is extremally charged to ensure the system is static.) Secondly, the black holes may have non-zero spin a , which could be a significant fraction of the Kerr bound $a = M$. Thirdly, a dynamical binary has orbital angular momentum, i.e., the two black holes orbit one another about the centre of mass, so p_ϕ will not be a constant of the motion. Lastly, the black holes will spiral inwards as they emit energy in the form of gravitational waves; the system will therefore not be static (or even stationary), so p_t will not be conserved.

Recall that Bohn *et al.* [190] have successfully performed ray-tracing to study strong-field gravitational lensing on a dynamical binary black hole spacetime. However, a range of other authors have adopted methods similar to those presented here, choosing instead to analyse closed-form models and surrogates. Let us briefly review some key studies here.

The Majumdar–Papapetrou di-hole spacetime is a subfamily of more general classes of exact solutions to the Einstein–Maxwell equations, including (i) the Israel–Wilson spacetimes [231]; (ii) the higher-dimensional Majumdar–Papapetrou solutions [217]; (iii) the Kastor–Traschen class of cosmological multi-black-hole spacetimes [200]; (iv) the Bretón–Manko–Aguilar solutions [232]; and (v) the double-Kerr–Newman solution [233].

In [198, 199] the authors study the Kastor–Traschen di-hole solution with a positive cosmological constant. This has the effect of pushing the black holes together, which mimics a head-on collision between the black holes. The shadows exhibit the same eyebrow-like features as those of the Majumdar–Papapetrou di-hole discussed in the present chapter.

A static solution to the Einstein field equations comprising a pair of uncharged black holes was first studied by Bach and Weyl [222]. In the absence of charges, a “Weyl strut” (i.e., a conical singularity along the symmetry axis connecting the sources) is required to maintain equilibrium between the black holes. This solution, occasionally referred to as the double-Schwarzschild black hole, was studied in the context of geodesic motion by Coelho and Herdeiro in 2009 [223].

More recently, Cunha *et al.* [234] have employed a novel method to reproduce the lensing effects of dynamical binaries (e.g. those studied by Bohn *et al.* [190]), by considering “quasi-static” ray-tracing on spacetimes given by closed-form solutions to the Einstein field equations, including the static double-Schwarzschild binary [222] and the stationary

double-Kerr binary [235]. In the latter case, the authors produce images of the shadows of the exact stationary co-rotating and counter-rotating stationary double-Kerr binary black hole configurations, which allows one to assess the impact of the intrinsic spin of the black holes on the binary shadows, contrasting it with the effect of the orbital angular momentum of the system. In related work [236], Cunha *et al.* use the double-Schwarzschild binary black hole to analyse whether black hole shadows can be used as a probe of the geometry of the black hole’s event horizon. The authors conclude that the conical deficit along the symmetry axis leaves no imprint on the black hole shadow.

Wang *et al.* [237] have studied the shadows cast by the Bonnor di-hole. They find that the presence of a magnetic dipole renders the null geodesic motion non-integrable, which results in chaotic patterns in the black hole shadows similar to those of the Majumdar–Papapetrou di-hole.

In this chapter, we have focussed on the physically interesting case of the shadow of a pair of black holes. Chaotic scattering of photons, however, is not unique to this scenario. In fact, assemblages of $N \geq 3$ black holes will exhibit even richer chaotic phenomena. Intriguingly, there are a number of singleton compact objects which have been shown to exhibit chaotic signatures, including perturbed or tidally distorted black holes [133, 238–240], Kerr black holes with scalar hair [191, 241] and boson stars [184, 242].

Wave propagation on binary black hole spacetimes

In the case of isolated black holes, (unstable) light-rings are related to important physical phenomena, such as strong gravitational lensing (e.g. the black hole shadow), the absorption cross-section of radiation incident on the black hole, and the relaxation of the black hole through quasinormal ringing. If two black holes are present (e.g. during the inspiral phase of a binary black hole merger), the non-escaping photon orbits exist, but their relationship to other aspects of binary black hole physics has not been well studied.

In this work, we have explored the relationship between non-escaping perpetual orbits of the Majumdar–Papapetrou binary black hole system and its black hole shadows. Recently, Assumpção *et al.* [225] have provided the first steps towards an understanding of wave dynamics around black hole binaries by considering the Majumdar–Papapetrou di-hole and the double-sink solution, a fluid mechanical analogue which describes a pair of “sonic holes”. In particular, the authors evolve the Klein–Gordon equation for a massless scalar field on the $(2 + 1)$ -dimensional double-sink geometry and explore the connection between the perturbations and the fundamental photon orbits (considered in this work). Further aspects of the physics of black hole binaries – including null geodesics, ringdown modes and energy extraction processes – have been studied by Bernard *et al.* [243].

4.9.2 Conclusions

Here, we outline the main conclusions of the work presented in this chapter, and discuss the possible implications of our results for gravitational lensing by dynamical binaries.

Firstly, we have demonstrated that chaotic scattering arises generically in the context of null geodesics motion around black holes when the spacetime admits more than one fun-

damental photon orbit, provided that such orbits are distinct and dynamically connected (i.e., a null geodesic may travel between the asymptotic neighbourhoods of the fundamental null orbits). The existence of more than one fundamental null orbit generates an uncountable infinity of non-escaping null orbits, which are neither absorbed nor scattered by the black holes. These orbits correspond to a fractal set of scattering singularities in the initial data (i.e., the strange repeller), for which the scattering process is undefined.

One key insight of this work is that black hole shadows can be viewed as the exit basins of an open Hamiltonian system. In the case of Majumdar–Papapetrou di-hole shadows, the (strange) repeller corresponds to the (fractal) boundary of the black hole shadows. One can therefore employ techniques from the theory of chaotic dynamical systems to understand chaotic scattering by binary black holes and the associated fractal structure of the black hole shadows.

Using decision dynamics for the Majumdar–Papapetrou di-hole (Section 4.5.3), we have been able to successfully describe the ordering of the non-escaping orbits in initial data in the case of null geodesic scattering in the meridian plane (Section 4.6.1). Moreover, we demonstrated that the one-dimensional binary black hole shadow may be constructed using an iterative procedure which involves deleting open intervals of initial data corresponding to scattered or absorbed rays; one may understand the procedure with the aid of symbolic dynamics (Section 4.6.2). This iterative process is similar to that which is used to construct the Cantor set. By considering a particular one-parameter scattering problem, we demonstrated that the one-dimensional shadow is manifestly self-similar (Section 4.6.3).

When motion is restricted to the meridian plane ($p_\phi = 0$), there exist three distinct fundamental periodic orbits (Figure 4.4). We find that these orbits are dynamically connected: a null ray can transition between the asymptotic neighbourhoods of these fundamental orbits by making a sequence of “decisions”, which are described by our choice of symbolic dynamics.

For $p_\phi \neq 0$, motion is not confined to the meridian plane. In this case, the existence of the fundamental photon orbits depends on the value of the angular momentum p_ϕ . We find that a two-dimensional shadow image can be decomposed into one-dimensional shadows, each of which corresponds to a value of $p_\phi = \text{constant}$. When transitions between fundamental orbits are not possible, the corresponding one-dimensional shadow is regular, rather than fractal.

Non-planar motion and the effect of angular momentum can be understood through the introduction of an effective potential (or height function) for null geodesics. It is favourable to work with this effective potential (rather than the geodesic potential) as it is independent of the photon’s orbital parameters, i.e., energy and angular momentum. The circular null orbits can be understood by classifying the stationary points of the effective potential.

Using the equal-mass Majumdar–Papapetrou di-hole as a case study, we analysed the effect of varying the coordinate separation parameter d . As the black holes are brought together, the morphology of the contours of the effective potential is affected. This changes the character of the fundamental periodic photon orbits. Once the value of d is sufficiently

small, light-rings about the individual components of the binary are forbidden; such orbits are supplanted by a fundamental periodic orbit which is symmetrical about the equatorial plane. Of course, a more detailed analysis would be required to understand the case of an unequal-mass binary.

Intriguingly, we uncover the possibility of *stable bounded photon orbits* around the two black holes in the Majumdar–Papapetrou di-hole spacetime for coordinate separations in the range $\sqrt{\frac{16}{27}} < \frac{d}{M} < \sqrt{\frac{32}{27}}$. The existence and phenomenology of stable photon orbits will be investigated in the context of stationary axisymmetric spacetimes in Chapter 6 of this thesis. In the case of quasi-bounded orbits, there exists a “pocket” feature in phase space with three narrow escapes which lead to the two black holes and spatial infinity. This acts as a “randomising region”, and yields qualitatively different chaotic behaviour which is associated with “highly fractalised” regions of the binary black hole shadow (see Figure 4.17). This will be explored further in Chapter 5.

We presented a practical method to track rays through event horizons in the maximally extended Majumdar–Papapetrou di-hole spacetime, based on the construction of ingoing/outgoing null coordinates centred on the event horizon of each black hole. When allowing for this possibility, we see that there is the potential for richer behaviour in the periodic orbits and chaotic orbits in the maximally extended spacetime.

Let us conclude with some comments on the physical implications of this work for real dynamical binary black holes. The Majumdar–Papapetrou di-hole, considered here as a surrogate model, is highly symmetric. The staticity and axisymmetry of the model mean that the momenta p_t and p_ϕ are conserved along null geodesics. Moreover, the two-dimensional shadow images can be decomposed into one-dimensional shadows, each of which corresponds to a fixed value of p_ϕ . In these one-dimensional shadows, we observed Cantor-like, regular, and “highly fractalised” structures, which could be understood by considering the existence of the fundamental photon orbits and the contours of the effective null geodesic potential. Of course, a dynamical binary will be neither static nor axially symmetric. It is therefore an open question whether all three types of structure are permitted in realistic binary black hole shadows. To investigate this possibility further, one may wish to analyse a sample of the ray-traced simulations generated by Bohn *et al.* [190], and investigate the existence of fundamental periodic null orbits in this case.

Chapter 5

Fractal structures in binary black hole shadows

5.1 Introduction

A key feature of general relativity is the gravitational lensing of light by massive bodies (e.g. stars or black holes). Such objects generate spacetime curvature, which results in the deviation of null geodesics on the curved background; see Chapter 3 for a review.

Black holes, which possess an event horizon, will cast a shadow – a two-dimensional region of an observer’s local sky which cannot be illuminated by distant light sources due to the blockage of the black hole. A shadow can equivalently be conceived as the set of all photon initial conditions on the observer’s “image plane” which, when traced backwards in time from the observer, asymptote towards the event horizon of the black hole. In the language of non-linear dynamics, black hole shadows can be viewed as exit basins (see Section 2.6.5) in an open Hamiltonian system with two escapes – one of which corresponds to the event horizon of the black hole; the other to spatial infinity.

For the scattering of null geodesics by a binary black hole system, such as the static Majumdar–Papapetrou di-hole (see Chapter 4), there exist three escapes: a photon may plunge into either of the black holes, or it may escape to spatial infinity. It is therefore natural to associate three exit basins with these escapes. In the case of the Majumdar–Papapetrou di-hole, both the exit basins in phase space and the black hole shadows exhibit a rich variety of structure: they may possess both fractal and regular (i.e., non-fractal) regions; moreover, we shall show that, in certain parameter regimes, the exit basins may possess the stronger *Wada property*, in which all three basins share a common fractal boundary.

In 1917, Yoneyama [244] proposed the *lakes of Wada* as a curious example of three open sets in \mathbb{R}^2 which all share the *same* fractal boundary; Yoneyama attributed the construction of the lakes of Wada to his supervisor Takeo Wada, after whom they are named. Close to the end of the century, in 1991, Kennedy and Yorke [245] demonstrated that the Wada property is not only a topological curiosity, but that three or more open sets which share a common fractal boundary may arise (generically) in non-linear dynamical systems. In the years that followed, the Wada property has been shown to exist in the

basins of a range of chaotic systems, including the Hénon–Heiles Hamiltonian system [117, 246], the Gaspard–Rice three-disc system [120, 226], and the Duffing oscillator [247]. For a comprehensive review, see [112].

A key consequence of the existence of Wada basins in the phase space of scattering problems is the difficulty which arises when predicting the final state of the system. If there is some small uncertainty in fixing the initial conditions close to a Wada boundary, this leads to a high level of indeterminacy and an extreme sensitive dependence on initial conditions, despite the system being completely deterministic. In the context of scattering by a pair of black holes, a photon whose initial conditions are chosen close to a Wada boundary in phase space could end up in one of three final states: the photon could plunge into either black hole or escape to spatial infinity.

In this chapter, we employ a recently developed numerical method [248] to test for the existence of Wada basins. This method relies on a simple observation: the boundaries of Wada basins are invariant under the pairwise merging of the basins. In [248], Daza *et al.* develop the “merging method”, and apply it to three canonical systems which are known to exhibit the Wada property: the forced damped pendulum; the Newton fractal; and the Hénon–Heiles Hamiltonian. The merging method takes as its input only the exit basin diagram at finite resolution (e.g. an image of the basins in phase space or the black hole shadows). Practically, the merging method determines whether an exit basin diagram possesses the Wada property up to a certain resolution.

The majority of the work presented in this chapter is based on [249]. In Section 5.2, we review the Majumdar–Papapetrou di-hole system – used here as a toy model for a more realistic dynamical black hole binary – from the perspective of Hamiltonian dynamics. We review the spacetime geometry and its geodesics (Section 5.2.1); describe the exit basins in phase space (Section 5.2.2); and describe the structure of the binary black hole shadows (Section 5.3). Section 5.4.1 contains a review of two quantitative measures of fractality in exit basins: the fractal dimension and the uncertainty exponent. In Section 5.4.2, we apply these methods to the Cantor basins as a pedagogical example, before using them to distinguish between regular and fractal parts of binary black hole shadows in Section 5.4.3. We then turn our attention to the more restrictive property of Wada. We review the Wada property in Section 5.5, with an emphasis on the exit basins of open Hamiltonian systems. In Section 5.6, we outline how one can verify that exit basins exhibit the Wada property using numerical algorithms. Section 5.6.1 contains a review of the merging method, which was first presented in [248]. We apply this method to the exit basins in phase space (Section 5.6.2), and to the binary black hole shadows of the Majumdar–Papapetrou di-hole (Section 5.6.3). We discuss our results and potential extensions to this work in Section 5.7.

5.2 Majumdar–Papapetrou open Hamiltonian system

The Majumdar–Papapetrou di-hole is a static axisymmetric solution to the Einstein–Maxwell equations, which describes the exterior spacetime of a pair of extremally charged Reissner–Nordström black holes, each with its mass parameter equal to its charge param-

eter ($M = Q$). This model was introduced and discussed in detail in Chapter 4. For ease of reference, we review the key features of the model here, from the perspective of Hamiltonian dynamics.

5.2.1 Spacetime geometry and Hamiltonian formalism for null rays

In Weyl–Lewis–Papapetrou coordinates $\{t, \rho, z, \phi\}$, the line element for the Majumdar–Papapetrou di-hole is given by (4.1) with one-form potential (4.2). Choosing $M_{\pm} = M$ (equal-mass case), the function appearing in the metric and electromagnetic potential takes the form

$$U(\rho, z) = 1 + \frac{M}{\sqrt{\rho^2 + (z - \frac{d}{2})^2}} + \frac{M}{\sqrt{\rho^2 + (z + \frac{d}{2})^2}}. \quad (5.1)$$

We recall that M is the mass of each black hole, and d is the separation between the centres. We hereafter employ units in which $M = 1$. In the chosen coordinate system, the black holes' event horizons appear as points, located on the symmetry axis at $z = \pm \frac{d}{2}$; we recall that these are actually null surfaces of topology $\mathbb{R} \times S^2$.

The geodesics of the Majumdar–Papapetrou di-hole $q^a(\lambda)$ are the integral curves of Hamilton's equations, with Hamiltonian $H(q, p) = \frac{1}{2}g^{ab}p_ap_b$, where g^{ab} are the contravariant components of the metric tensor, $p_a = g_{ab}\dot{q}^a$ are the canonical momenta, and an overdot denotes differentiation with respect to the affine parameter λ . The time independence and axial symmetry of the spacetime mean that t and ϕ are ignorable coordinates; thus, their conjugate momenta p_t and p_ϕ are constants of motion. Moreover, the Hamiltonian H is conserved along geodesics. In the case of null geodesics, we have $H = 0$, and we may set $p_t = -1$ without loss of generality, by availing the freedom to rescale the affine parameter.

We recall that, after performing a conformal transformation of the metric, we are able to express the Hamiltonian as

$$H = \frac{1}{2}(p_\rho^2 + p_z^2) + V, \quad V(\rho, z) = -\frac{1}{2\rho^2}(h - p_\phi)(h + p_\phi), \quad (5.2)$$

where the *effective potential* (or *height function*) is

$$h(\rho, z) = \rho U^2, \quad (5.3)$$

and p_ϕ is a free parameter which is conserved along rays (see Section 4.3.2).

The null condition $H = 0$ and the positivity of the kinetic term in the Hamiltonian (5.2) together imply that the geodesic potential must satisfy $V(\rho, z) \leq 0$; this inequality defines the *allowed regions* in configuration space. The contours $h = \pm p_\phi$ (which are equivalent to $V = 0$) demarcate the boundary of the allowed regions in the (ρ, z) -plane.

The full phase space, spanned by the four spacetime coordinates and their conjugate momenta, $\{q^a, p_b\}$, is eight-dimensional. The symmetries of the spacetime give rise to conserved quantities (associated with the ignorable coordinates t and ϕ). This permits us to focus on a four-dimensional *reduced* phase space, with two pairs of conjugate variables $\{\rho, z, p_\rho, p_z\}$, and an additional constraint $H = 0$. The null condition $H = 0$ allows one of the coordinates in the reduced phase space (e.g. p_z) to be expressed as a function

of the other three coordinates: the motion is restricted to a three-dimensional energy hypersurface which lies in the reduced phase space, described by $H = 0$.

In the reduced phase space, the Majumdar–Papapetrou di-hole shares qualitative features with the Hénon–Heiles Hamiltonian system [117], which has become a paradigm for two-dimensional time-independent Hamiltonian scattering since its introduction in the 1960s. We review the Hénon–Heiles system in Appendix C.

5.2.2 Exit basins in phase space

Exit basins of the Majumdar–Papapetrou di-hole

Here, we construct exit basins in a two-dimensional subspace of the reduced phase space for the Majumdar–Papapetrou di-hole with $d = 1$. The initial conditions used to construct the exit basins are chosen in close analogy with those of a study of the exit basins of the Hénon–Heiles Hamiltonian system [246].

To construct the exit basins, we focus on the case of equal-mass black holes separated by coordinate distance $d = 1$, due to the high symmetry of this scenario. The effective potential h possesses three saddle points, one of which is in the equatorial plane at $(\rho, z) = (\frac{1}{2}5^{1/4}\varphi^{3/2}, 0)$, and the other two are out of the plane at $(\rho, z) = (\frac{1}{2}5^{1/4}\varphi^{-1/2}, \pm\frac{1}{2\varphi})$, where $\varphi = \frac{1}{2}(1 + \sqrt{5})$ is the golden ratio. All three of the saddle points lie on a single critical contour $h = p_\phi^* = \frac{1}{2}5^{5/4}\varphi^{3/2}$. The critical contour encloses a local maximum of h , which is located at $(\rho, z) = (\rho_{(1)}, z_{(1)}) = (\frac{\sqrt{3}}{2}, 0)$, with $h = p_\phi^{(1)} = \frac{9\sqrt{3}}{2}$. A derivation of the closed-form expressions for the location of the fixed points and critical contours of h is given in Appendix B. The stationary points and contours of h are presented in Figure 5.1(a). Contrasting Figures 5.1(a) and C.1, one can clearly see a qualitative resemblance between the contours of the Majumdar–Papapetrou effective potential (5.3) and those of the Hénon–Heiles potential (C.2).

For $p_\phi \geq p_\phi^*$, the equipotential curves $h = p_\phi$ form closed curves on a subregion of the (ρ, z) -plane, close to the local maximum of h : there is a closed region of phase space, in which orbits are kinematically bounded. Conversely, for $p_\phi < p_\phi^*$, the contours $h = p_\phi$ are open: null rays are permitted to plunge into either of the black holes or escape to infinity. In this chapter, we focus on the latter case, treating the Majumdar–Papapetrou di-hole as a novel example of a two-dimensional time-independent open Hamiltonian system with three escapes in phase space. We express the azimuthal angular momentum as $p_\phi = p_\phi^* - \Delta p_\phi$, where $0 < \Delta p_\phi \leq p_\phi^*$.

Our first choice of initial conditions is to fix the coordinates ρ and z , and choose the initial two-momentum (p_ρ, p_z) to be tangent (in the anti-clockwise sense) to the circle of radius $\sqrt{(\rho - \rho_{(1)})^2 + (z - z_{(1)})^2}$, centred on the maximum of h , which is located at $(\rho_{(1)}, z_{(1)}) = (\frac{\sqrt{3}}{2}, 0)$; see Figure 5.1(b) for the set-up. The corresponding exit basins are plotted in the (ρ, z) -plane; see Figures 5.2(a)–5.2(c) for the exit basin diagrams. The second choice of initial conditions is to fix $z = 0$, and vary the values of ρ and p_ρ , as shown in Figure 5.1(c). The exit basin diagrams are plotted in the (ρ, p_ρ) -plane; see Figures 5.2(d)–5.2(f).

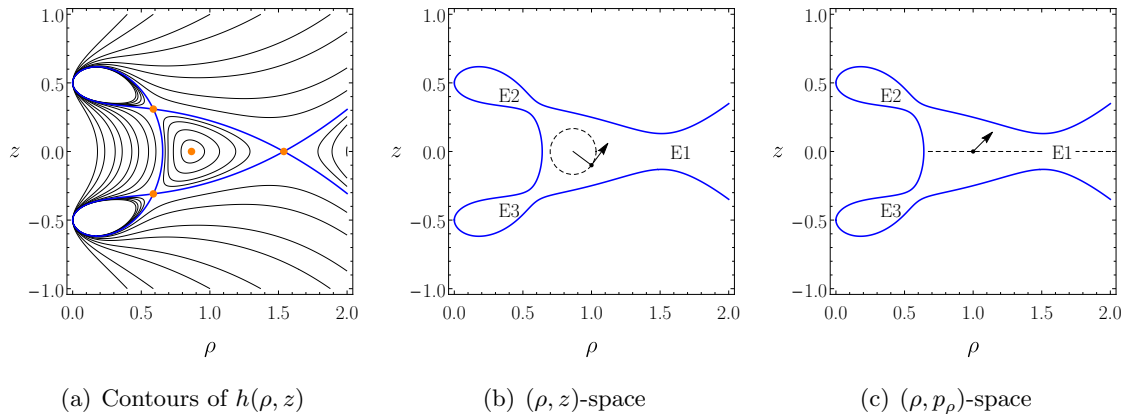


Figure 5.1. (a) Contours of the effective potential $h(\rho, z) = \rho U^2$ for the equal mass ($M_\pm = M = 1$) Majumdar–Papapetrou di-hole with separation parameter $d = M = 1$. A critical contour passes through three saddle points, and encloses a local maximum. (b) Initial conditions for exit basins in the (ρ, z) -plane. The initial two-momentum (p_ρ, p_z) is tangent to the circle centred on the local maximum of h that passes through the initial condition. (c) Initial conditions for exit basins in the (ρ, p_ρ) -plane. The photon is fired from the ρ -axis ($z = 0$) and the initial values of ρ and p_ρ are varied.

In Figure 5.2, we present a selection of exit basins for various values of the angular momentum parameter $\Delta p_\phi = p_\phi^* - p_\phi$, employing both choices of initial conditions described above. To realise the exit basins in phase space, we consider a fine grid of initial conditions in either the (ρ, z) -plane or the (ρ, p_ρ) -plane. We then integrate Hamilton’s equations subject to these initial conditions, until a halting condition is satisfied: the numerical integrator is stopped when either $\rho^2(\lambda) + (z(\lambda) - z_\pm)^2 \leq \varepsilon$, where $0 < \varepsilon \ll 1$; or when $\rho^2(\lambda) + z^2(\lambda) \geq R$, where $R \gg 1$ is some sufficiently large radius. The former case corresponds to absorption by the black hole located at $z = z_\pm$, and the latter case indicates scattering.

In order to visualise the exit basins, we colour the initial conditions grey if they lead to infinity, green for the upper black hole, and purple for the lower black hole (cf. the shadows of Chapter 4). In each of the exit basins diagrams of Figure 5.2, we observe an intricate fractal-like structure. For larger values of Δp_ϕ , the basins are more regular with “sharper” boundaries. As $\Delta p_\phi \rightarrow 0$, the three basins become more intertwined and exhibit “irregular” structure. Moreover, in this limit, Kolmogorov–Arnold–Moser (KAM) islands of stability [112] (depicted in black) appear. The KAM islands consist of the initial conditions corresponding to trajectories which never escape the scattering region as $\lambda \rightarrow \infty$, despite the system being open (i.e., $\Delta p_\phi > 0$). Trajectories inside KAM islands never escape to infinity nor end up in either of the black holes; rather, they wander forever through the scattering region in a quasi-periodic fashion. The KAM islands are organised in a fractal hierarchy, and have non-zero (Lebesgue) measure, as can be inferred from Figures 5.2(c) and 5.2(f). As the value of Δp_ϕ is increased (i.e., the width of the escape channels increases), all trajectories escape the scattering region through one of the three escapes, and the KAM islands disappear. A similar effect is seen in the Hénon–

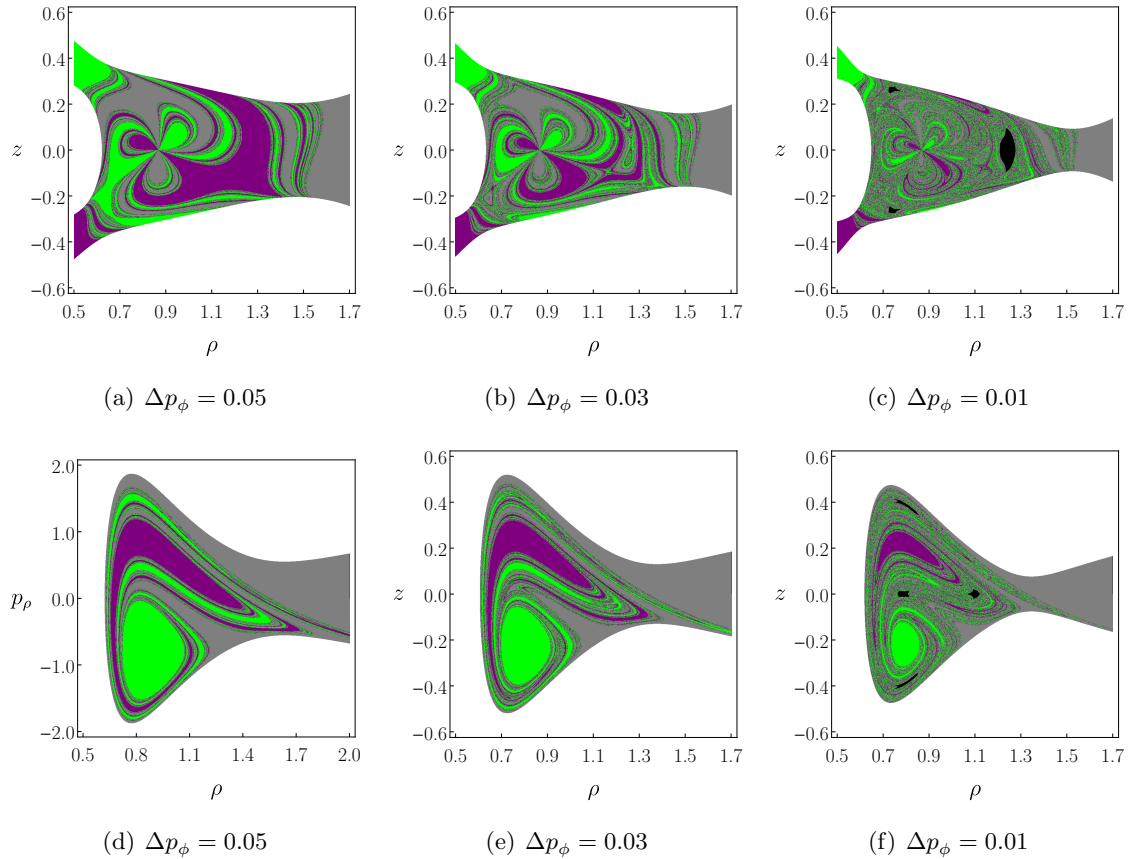


Figure 5.2. Exit basins in the Majumdar–Papapetrou open Hamiltonian system for (a)–(c) initial conditions in (ρ, z) -space; and (d)–(f) initial conditions in (ρ, p_ρ) -space. (Both choices of initial conditions are inspired by studies of the Hénon–Heiles system, and are described in the text.) The initial data which lead to spatial infinity are plotted in grey, and those which asymptote to the event horizon of the upper (lower) black hole are shown in green (purple). As one decreases the parameter $\Delta p_\phi = p_\phi^* - p_\phi$, the widths of the three escapes decrease, and KAM islands of stability (depicted in black) appear.

Heiles Hamiltonian system as one increases the total energy E . A general discussion of the limit of small escapes in open Hamiltonian systems is presented in [227], where the authors employ the Gaspard–Rice three-disc model and the Hénon–Heiles Hamiltonian as illustrative examples.

Comparison with the Hénon–Heiles Hamiltonian system

The Hénon–Heiles system was investigated in the context of chaotic scattering in [246], for conserved energies in the range where trajectories are permitted to escape from the scattering region (see Appendix C for a summary). Exit basins were employed to develop a qualitative understanding of the dynamics in the (open) Hénon–Heiles Hamiltonian system. The authors demonstrate, using various computational methods, that (in a certain parameter regime) the exit basins are not only fractal, but possess the more restrictive Wada property; that is, each point on a basin boundary belongs to the boundary of all

three basins.

By comparing the exit basin diagrams, one can clearly see a resemblance between the Majumdar–Papapetrou di-hole system (Figure 5.2) and the Hénon–Heiles system (Figure C.3). The basins appear to be “topologically equivalent”; however, the Hénon–Heiles basins are more symmetric due to the rotational symmetry of the potential. Given the striking similarities between the exit basins of two systems, and that the basins of the Hénon–Heiles system are known to exhibit the Wada property, we anticipate that the basins of the Majumdar–Papapetrou system will satisfy the same property. This will be explored in more depth in Section 5.6.

5.3 Transition in shadow structure

5.3.1 Structure of Majumdar–Papapetrou di-hole shadows

Here, we consider shadows of the equal-mass ($M_{\pm} = M = 1$) Majumdar–Papapetrou di-hole, where the separation between the centres is d . The shadows were constructed using ray-tracing in Section 4.7.2. In Figures 4.15–4.17, we present Majumdar–Papapetrou di-hole shadows on the observer’s image plane – spanned by the coordinates (X, Y) – for a range of viewing angles θ , and separations d .

In the equatorial case ($\theta = \frac{\pi}{2}$), the two-dimensional shadow images can be viewed as the union of one-dimensional shadows of constant p_{ϕ} (equivalently Y); see Figure 4.17. In Section 4.7.2, we saw that the one-dimensional shadows can either exhibit Cantor-like (fractal) structure, or they can be regular. As a consequence, the two-dimensional binary black hole shadow image can either be “homogeneous” or “heterogeneous”. In the former case, each one-dimensional shadow is Cantor-like; in the latter, the two-dimensional shadow is the union of Cantor-like shadows and regular shadows. In particular, we saw that the shadow for $d = 1$ is homogeneous (Figures 4.17(d)–4.17(f)), whereas the shadow for $d = 2$ is heterogeneous (Figures 4.17(a)–4.17(c)).

We anticipate that there will be some critical value of the separation, say $d = \hat{d}$, such that two-dimensional Majumdar–Papapetrou di-hole shadows with $d < \hat{d}$ will be homogeneous, whereas shadows with $d > \hat{d}$ will be heterogeneous. In Section 5.3.2, we describe the role of the so-called fundamental photon orbits in the structure of one-dimensional shadows. Then, in Section 5.3.3, we provide a practical method which can be used to calculate the critical value of the separation parameter \hat{d} , which marks the “phase transition” in the shadow structure from homogeneous to heterogeneous.

5.3.2 Role of photon orbits

Expanding on the discussion in Sections 4.7.2 and 5.3.1, we now turn our attention to the so-called fundamental photon orbits, a special class of null rays which play an important role in the Majumdar–Papapetrou di-hole spacetime. We recall from Section 5.2 that the allowed regions of phase space, accessible to a null geodesic with azimuthal angular momentum p_{ϕ} , are demarcated by the contours of the effective potential (5.3). For an equal-mass Majumdar–Papapetrou di-hole, a *fundamental photon orbit* is a null geodesic

$q^a(\lambda)$ which satisfies the following properties: (i) it is restricted to a compact subset of the (ρ, z) -plane; (ii) it is periodic, i.e., there exists some $T > 0$ such that $q^a(\lambda) = q^a(\lambda + T)$ for all $\lambda \in \mathbb{R}$; (iii) it is unstable; (iv) it touches the contour $h(\rho, z) = p_\phi$ in such a way that, locally, it is orthogonal to the contour; and (v) the radial momentum p_ρ vanishes if the orbit passes through the equatorial plane, by symmetry.

More general null geodesics are permitted in the Majumdar–Papapetrou di-hole space-time which satisfy some (but not all) of properties (i)–(v). For example, null geodesics which satisfy (i) and (ii) but which are *stable* are permitted; these orbits are explored in Chapter 6. In [79], Cunha *et al.* present a classification scheme for generic fundamental photon orbits in stationary axisymmetric spacetimes, which need only satisfy properties (i) and (ii). The role of light-rings and generic fundamental photon orbits in the analysis of strong-field gravitational lensing is discussed in [60]. (In this work, we reserve the label “fundamental” for a photon orbit which satisfies all of the properties (i)–(v) from the above list.)

Figure 4.11 shows examples of the three types of fundamental photon orbit around the equal-mass Majumdar–Papapetrou di-hole with $d = 2$, projected onto the (ρ, z) -plane. These fundamental orbits were discussed in the case of motion restricted to the meridian plane ($p_\phi = 0$) and non-planar motion ($p_\phi \neq 0$) in Chapter 4. Moreover, these orbits can be understood using symbolic dynamics (see Section 4.5.3). For the remainder of this chapter, we label the three types of fundamental photon orbits as follows: (I) a one-component light-ring around only one of the black holes; (II) a figure-of-eight orbit around both black holes; (III) a two-component light-ring around both black holes. Recall that, in decision dynamics (Section 4.5.3), the fundamental orbits have representations (I) $\bar{0}$; (II) $\bar{2}$; (III) $\bar{4}$.

In Figures 4.11(a)–4.11(d), we demonstrate the effect of increasing p_ϕ on the fundamental orbits, for the case of equal-mass black holes separated by coordinate distance $d = 2$. As one increases p_ϕ from zero, the contour $h(\rho, z) = p_\phi$ moves away from the symmetry axis and orbits II and III move closer together. These orbits merge together at $p_\phi = \hat{p}_\phi \approx 5.09$; beyond this value ($p_\phi > \hat{p}_\phi$), the type II and III orbits no longer exist. We see that the type I orbits persist until $p_\phi = p_\phi^*$, which is the value of the azimuthal angular momentum for the pair of non-planar unstable circular photon orbits of constant ρ and z . These unstable circular orbits correspond to the non-planar saddle points of the effective potential $h(\rho, z)$; see Figure 4.11(d). We note that a ray with $p_\phi > p_\phi^*$ which starts far from the scattering region is forbidden from accessing the black holes by the contour $h = p_\phi$; the value p_ϕ^* corresponds to the top of the binary black hole shadow image (i.e., the last value of Y for which absorption is permitted); see Figure 4.17.

As demonstrated in Chapter 4, when all three types of orbit exist and are “dynamically connected” – i.e., if rays are permitted to transition between the asymptotic neighbourhoods of two or more fundamental orbits – then chaotic scattering arises naturally. Moreover, we find that this chaotic scattering is responsible for Cantor-like fractal structures in one-dimensional binary black hole shadows with $p_\phi = \text{constant}$. For a given value of d , the fundamental photon orbits are dynamically connected for $p_\phi < \hat{p}_\phi$. Indeed, in Chapter 4, we demonstrated that, for a given value of $p_\phi < \hat{p}_\phi$, the one-dimensional shadow has a

Cantor-like structure; see Figure 4.17. We note that it is not sufficient for two separate orbits of type I to exist: these are not typically dynamically connected in the absence of type II and III orbits.

In Chapter 4, we describe how the structure of the one-dimensional shadows for coordinate separation $d = 2$ changes as p_ϕ varies. For $p_\phi > p_\phi^*$, the ray is forbidden from accessing the black holes by the contour $h = p_\phi$: the corresponding one-dimensional shadow is the empty set. As describe above, when $\hat{p}_\phi < p_\phi < p_\phi^*$, the orbits of type II and III do not exist and the type I orbits are not dynamically connected: the shadow is therefore regular in this regime. Finally, for $0 < p_\phi < \hat{p}_\phi$, all orbits of types I–III exist and are dynamically connected: the one-dimensional shadow exhibits a Cantor-like fractal structure.

Heterogeneous shadows are admitted by equal-mass di-holes with coordinate separation d , such that the coexistence condition $\hat{p}_\phi < p_\phi^*$ is satisfied. In the following section, we demonstrate that this condition is only met for sufficiently separated black holes with $d > \hat{d}$; for $d < \hat{d}$, the coexistence condition is not met, and the two-dimensional shadow will be homogeneous (i.e., each one-dimensional shadow will be Cantor like).

5.3.3 Calculating the critical separation

Here, we briefly outline a method to calculate the critical value of the separation \hat{d} , introduced in Section 5.3.2. The critical separation $d = \hat{d}$ marks a transition from homogeneous ($d < \hat{d}$) to heterogeneous ($d > \hat{d}$) binary black hole shadows.

Recall from Section 5.3.2 that, for a fixed value of d , the type II and III fundamental photon orbits merge at $p_\phi = \hat{p}_\phi$. We seek the value of the coordinate separation between a pair of equal-mass black holes in the Majumdar–Papapetrou di-hole which gives rise to a single outer fundamental photon orbit when $\hat{p}_\phi = p_\phi^*$; that is, the value of d for which the type II and III fundamental photon orbits merge at exactly the value of p_ϕ for which the two black holes become inaccessible to a ray incident from infinity.

We first choose a value of the coordinate separation d . The location of the stationary points of $h(\rho, z)$, and thus the corresponding value of p_ϕ^* , are found using the method presented in Appendix B. We then consider rays which start on the contour $h = p_\phi^*$, with $\rho(0) = \rho_0$. The initial value $z(0) = z_0 > 0$ is found by numerically solving the equation $h(\rho_0, z_0) = p_\phi^*$ for z_0 . A ray which starts on the contour has $p_\rho(0) = 0 = p_z(0)$ from the Hamiltonian constraint $H = 0$.

Having fixed the initial conditions $(\rho(0), z(0), p_\rho(0), p_z(0)) = (\rho_0, z_0, 0, 0)$, we evolve the geodesic equations (i.e., Hamilton’s equations) until the ray passes through the equatorial plane ($z = 0$). At this point, we record the value of $\vartheta = \frac{\pi}{2} + \arctan\left(\frac{p_z}{p_\rho}\right)$, where $\arctan\left(\frac{p_z}{p_\rho}\right)$ is the angle made by the tangent vector (i.e., the two-momentum) and the equatorial plane. By symmetry, the type II and III fundamental orbits must have $\vartheta = 0$. Hence, the zeros of the function $\vartheta(\rho_0)$ determine the location of the initial conditions $\rho(0) = \rho_0$ (on the contour $h = p_\phi^*$) of the type II and III fundamental orbits.

In Figure 5.3, we show the function $\vartheta(\rho_0)$ for three representative values of the separation d . In each case, the curve has been determined numerically using the method described above. We are interested in the value $d = \hat{d}$ for which $\vartheta(\rho_0)$ admits a single

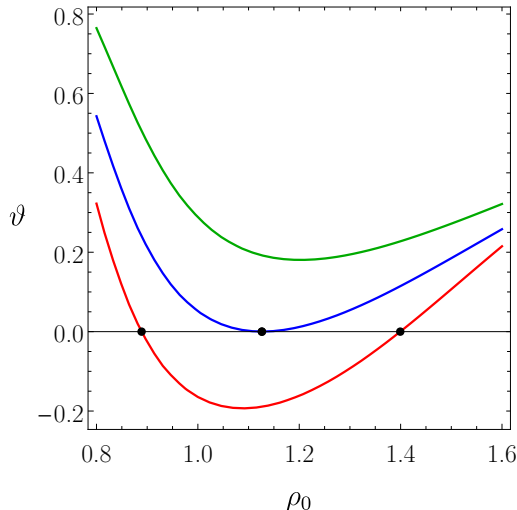


Figure 5.3. Results of the algorithm used to search for initial data $\rho(0) = \rho_0$ corresponding to the type II and III fundamental orbits, for which $\vartheta = 0$. We show three representative examples for varying coordinate separations d : (i) $d < \hat{d}$ [red]; (ii) $d = \hat{d}$ [blue]; (iii) $d > \hat{d}$ [green]. The function $\vartheta(\rho_0)$ has (i) two zeros, so both type II and III orbits exist for all $p_\phi \in [0, p_\phi^*]$; (ii) a single zero, so there is a single outer (type II/III) orbit; (iii) no zeros, so there exists some $\hat{p}_\phi < p_\phi^*$ beyond which the type II and III orbits no longer exist.

zero, i.e., there exists a single outer (type II/III) fundamental orbit. We find that

$$\hat{d} \approx 1.2085M \quad (5.4)$$

for the equal-mass Majumdar–Papapetrou di-hole. This is shown as a blue curve in Figure 5.3. In the case of unequal masses, the \mathbb{Z}_2 reflection symmetry is broken; a more detailed analysis of the structure of the fundamental orbits would be required to determine the value of \hat{d} in this case.

For di-holes with a sufficiently small separation ($d < \hat{d}$), there exist three types of fundamental orbits for all $p_\phi \in [0, p_\phi^*]$. In this regime, we anticipate that the boundary of the Majumdar–Papapetrou binary black hole shadow will be entirely fractal (homogeneous). On the other hand, for sufficiently separated black holes ($d > \hat{d}$), there exists some range of values $p_\phi \in (\hat{p}_\phi, p_\phi^*)$ in which the type II and III fundamental orbits no longer exist, but absorption by the black holes is still possible. (This phenomenon is demonstrated in Figure 4.11, where we consider the case $d = 2M > \hat{d}$.) In such situations, we anticipate that the regions of the binary black hole shadow boundary corresponding to $p_\phi \in (\hat{p}_\phi, p_\phi^*)$ will no longer be fractal. The two-dimensional shadow image will therefore be heterogeneous.

5.4 Uncertainty exponent and fractal dimension

5.4.1 Definitions and numerical method

Hausdorff dimension of fractal sets

Informally, a geometrical figure which contains m copies of itself which are scaled by a factor $\frac{1}{s}$ is said to have *Hausdorff dimension* D_H , where $m = s^{D_H}$ [99, 114]. (Note that

the Hausdorff dimension is often referred to as the *Hausdorff–Besicovitch dimension*.) Inverting this scaling law yields

$$D_{\text{H}} = \frac{\ln m}{\ln s}. \quad (5.5)$$

A more formal definition of the Hausdorff dimension, which relies on the notion of Hausdorff measure, is given in e.g. [114]. The definition (5.5) will be sufficient for the purpose of this work.

To illustrate the (heuristic) calculation of the Hausdorff dimension, let us consider some simple geometric objects. A line segment contains m equal-sized copies of itself, each of which is scaled by a factor $\frac{1}{m}$, where $m \in \mathbb{N}$. The Hausdorff dimension of a line segment is therefore $D_{\text{H}} = \frac{\ln m}{\ln m} = 1$. Similarly, a square can be divided into k^2 smaller squares, which are scaled by a factor $\frac{1}{k}$, where $k \in \mathbb{N}$. The Hausdorff dimension of a square is therefore $D_{\text{H}} = \frac{\ln k^2}{\ln k} = 2$. In both of these cases, the Hausdorff dimension takes an integer value which agrees with our usual notion of dimension (i.e., the topological dimension) [114].

The Hausdorff dimension is intimately related to the notions of scaling and self-similarity, and provides a natural way to characterise fractals. The Hausdorff dimension is therefore often referred to as the *fractal dimension*. In fact, some attempts to provide a rigorous mathematical definition of a fractal are dependent on the concept of Hausdorff dimension. For example, Mandelbrot [250] conceived that a fractal is a set whose Hausdorff dimension is strictly greater than its topological dimension. This definition proved to be unsatisfactory, as it excludes a large number of sets which clearly ought to be defined as fractals; for example, certain space-filling curves (e.g. the Hilbert curve) do not meet Mandelbrot’s criterion.

Consider the Cantor set, which is an example of a self-similar fractal constructed on the unit interval (see Section 4.6.2). Heuristically, we may calculate the Hausdorff dimension D_{H} of the Cantor set \mathcal{C} as follows. The Cantor set contains two copies of itself, $\mathcal{C} \cap [0, \frac{1}{3}]$ and $\mathcal{C} \cap [\frac{2}{3}, 1]$, each of which has been scaled by a factor $\frac{1}{3}$. The Hausdorff dimension of \mathcal{C} is thus $D_{\text{H}} = \frac{\ln 2}{\ln 3} \approx 0.631$. The Hausdorff dimension of the Cantor set is between zero (the Hausdorff dimension of a countable set) and one (the Hausdorff dimension of the unit interval).

Similarly, we may consider the γ -Cantor set \mathcal{C}_{γ} , defined in Section 4.6.2. This set contains two copies of itself which have been scaled by a factor $\frac{1-\gamma}{2}$. The Hausdorff dimension of \mathcal{C}_{γ} is therefore

$$D_{\text{H}}(\gamma) = \frac{\ln 2}{\ln \left(\frac{2}{1-\gamma} \right)}. \quad (5.6)$$

Varying $\gamma \in (0, 1)$ gives any Hausdorff dimension $D_{\text{H}}(\gamma) \in (0, 1)$. Clearly, the case $\gamma = \frac{1}{3}$ gives the Hausdorff dimension of the canonical middle-thirds Cantor set.

Uncertainty exponent

Grebogi *et al.* [251] considered the *uncertainty* (or *unpredictability*) associated with basin boundaries in dynamical systems, introducing the *uncertainty exponent* as a quantitative measure of the indeterminacy in final-state prediction when fractal boundaries are present.

Here, we review the definition of the uncertainty exponent, and discuss its usefulness in distinguishing between regular and fractal boundaries.

Consider an initial uncertainty ε in the determination of the initial conditions, which can, in general, be due to numerical or experimental error, or a small perturbation of the initial data. Given this uncertainty, we compute the fraction of initial conditions f which are *uncertain*, i.e., the proportion of initial conditions that are within a distance ε of the basin boundary. Computing f for different values of ε , we see that $f \sim \varepsilon$ for smooth boundaries. However, for fractal boundaries, the fraction of uncertain initial conditions obeys a power-law scaling of the form $f \sim \varepsilon^\alpha$. Here, α is the *uncertainty exponent*, which takes values between zero (total fractality) and one (regularity).

The uncertainty exponent is related to the dimension D of the space in which the object is embedded, and the Hausdorff dimension of the fractal boundary D_H , via [251]

$$\alpha = D - D_H. \quad (5.7)$$

(A proof of this statement is given in [252].) The uncertainty exponent of the Cantor basins, whose boundary is the middle-thirds Cantor set, is $\alpha = 1 - \frac{\ln 2}{\ln 3} \approx 0.369$. This indicates that the Cantor basins are indeed fractal.

Numerical method to calculate the uncertainty exponent

In practice, the uncertainty exponent is typically computed numerically by analysing how the fraction of uncertain initial conditions f varies with the uncertainty ε . For a given initial condition x , one first finds the basin to which x belongs. Next, one determines the basin for (a sample of) initial conditions y which satisfy $|x - y| \leq \varepsilon$, for some suitable distance measure $|\cdot|$. (Here, we use the Euclidean distance on \mathbb{R}^n .) If all of these initial conditions y belong to the *same* basin as x , then the point x is labelled *certain*, otherwise it is labelled *uncertain*.

For a fixed value of ε , a large number N of initial conditions x are chosen – either randomly or from a uniform grid. These are labelled as certain or uncertain according to the method described above. The total number of uncertain initial conditions is denoted N_u . Then $f(\varepsilon) = \frac{N_u}{N}$, the fraction of initial conditions which are uncertain. This procedure is repeated for a range of values of the tolerance parameter ε , in order to determine the power-law relationship $f \sim \varepsilon^\alpha$. The slope of the graph of $\log f$ versus $\log \varepsilon$ is precisely the uncertainty exponent α .

5.4.2 Uncertainty exponent for Cantor basins

To test and calibrate the method outlined in Section 5.4.1, we compute α numerically for the two-colour γ -Cantor basins $\{B_1, B_2\}$ (defined in Section 4.6.2), and we compare this with an exact result.

The numerical calculation of the uncertainty exponent is performed using a fixed set of $N = 10^8$ equally spaced “initial conditions” in the unit interval $I = [0, 1]$. For a fixed value of ε , we consider each of the N initial conditions $x \in I$ in turn. The initial condition x is classified as certain or uncertain using the procedure outlined above. That is, by looking

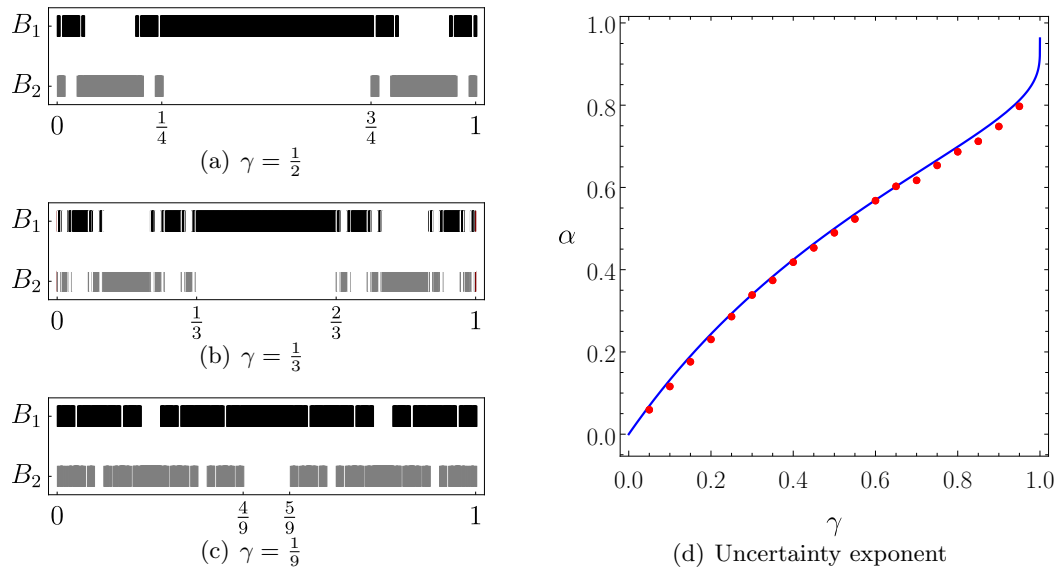


Figure 5.4. Uncertainty exponent for Cantor basins. (a)–(c) Two-colour γ -Cantor basins $\{B_1, B_2\}$ for a selection of values of γ . (d) Uncertainty exponent α for the γ -Cantor basins. The graph of the analytical expression is shown in blue; the numerical values, computed for $\gamma \in \{0.05, 0.10, \dots, 0.95\}$ using the method described in the text, are shown in red. The uncertainty exponent increases with γ . This indicates that the smaller γ is, the more uncertain (i.e., the more “fractalised”) the γ -Cantor basins are.

at the subset of the N initial conditions which are within a distance ε , and determining to which basin each one belongs. If all such points belong to the same basin as the initial condition under consideration, then the initial condition x is labelled as certain; otherwise, x is labelled uncertain. The uncertainty exponent α is then obtained from the slope of $\log f$ versus $\log \varepsilon$, where f is the fraction of uncertain initial conditions.

For the γ -Cantor basins, the uncertainty exponent $\alpha(\gamma)$ may be obtained in closed form by virtue of (5.6) and (5.7). That is, by subtracting the Hausdorff dimension of the γ -Cantor set – the boundary of the γ -Cantor basins – from unity (the topological dimension of the unit interval). This yields

$$\alpha(\gamma) = 1 + \frac{\ln 2}{\ln \left(\frac{1-\gamma}{2} \right)}, \quad \gamma \in (0, 1). \quad (5.8)$$

The function $\alpha(\gamma)$ is monotonically increasing on $\gamma \in (0, 1)$ with $\lim_{\gamma \rightarrow 0^+} \alpha(\gamma) = 0$ and $\lim_{\gamma \rightarrow 1^-} \alpha(\gamma) = 1$. This indicates that the basins become *more uncertain* (i.e., α increases) as we decrease the parameter γ .

In Figures 5.4(a)–5.4(c) we show examples of two-colour γ -Cantor basins for $\gamma \in \{\frac{1}{2}, \frac{1}{3}, \frac{1}{9}\}$. We consider the case of two basins because the uncertainty exponent is concerned only with the *boundary* of the basins, not the number of basins. We recall that the boundary of the γ -Cantor basins is the middle- γ Cantor set, no matter how many basins we choose to construct; see Section 4.6.2. One can see by inspection of Figures 5.4(a)–5.4(c) that the basins become more “fractalised” (i.e., more uncertain) as one increases γ .

Figure 5.4(d) shows the uncertainty exponent for the γ -Cantor basins, calculated using (i) the closed-form expression for the uncertainty exponent (5.8) [blue curve]; and (ii) the numerical method outlined above with $\gamma \in \{0.05, 0.10, \dots, 0.95\}$ and a resolution of $N = 10^8$ initial conditions [red points]. We see that there is good agreement between the analytical and numerical results. Better agreement could be obtained for the more uncertain basins (which have $\gamma \approx 1$) by calculating α numerically using a finer grid of initial conditions; however, this would be more computationally expensive.

5.4.3 Uncertainty exponent for binary black hole shadows

In this section, we compute the uncertainty exponent of the boundary of one-dimensional equal-mass Majumdar–Papapetrou di-hole shadows of constant Y , in the cases $d = 1$ and $d = 2$, for a viewing angle of $\theta = \frac{\pi}{2}$. Recall that each slice of constant Y is related to a fixed value of the azimuthal angular momentum p_ϕ , as described in Section 4.7.2. Here, we compute the uncertainty exponent using the numerical method outlined in Section 5.4.1, for one-dimensional shadows with $Y \in \{-8.0, 7.9, \dots, 8.0\}$, using $N = 10^7$ initial conditions in each slice. Once the uncertainty exponent has been calculated, one can obtain the Hausdorff dimension of the boundary using (5.7) with $D = 1$. In the following analysis, we discuss only $Y \geq 0$, because the equatorial ($\theta = \frac{\pi}{2}$) shadows are symmetric about $Y = 0$.

Let us first consider the case $d = 2$, shown in Figures 5.5(a) and 5.5(b). As outlined in Section 4.7.2, we observe Cantor-like and regular boundaries in the one-dimensional cross sections of this shadow; see Figure 4.17. We anticipate that the transition from Cantor-like to regular will occur when the outer (type II/III) fundamental orbits coincide. This happens when $p_\phi = \hat{p}_\phi \approx 5.09$, which corresponds to $Y \approx 4.71$. Figure 5.5(a) shows the uncertainty exponent for this case. This figure is transposed, so that Y runs along the vertical axis, to aid comparison with the shadow presented in Figure 5.5(b). The transition from regular to Cantor-like is indicated using a red horizontal line in Figures 5.5(a) and 5.5(b). We find that $\alpha < 1$ for $Y < 4.71$, so this part of the shadow is fractal; whereas $\alpha \approx 1$ for $Y > 4.71$, which indicates that this part of the shadow is not fractal. This agrees with our understanding of the shadow structure and the fundamental orbits, outlined in Chapter 4 and Section 5.3. Crucially, the uncertainty exponent is able to distinguish between the fractal and non-fractal parts of the shadow.

We now consider the uncertainty exponent for a di-hole with $d = 1$. Figure 5.5(c) shows that, in this case, $\alpha < 1$ across the whole shadow, which indicates that the entire shadow boundary is fractal. We observe that the uncertainty exponent attains its minimum value ($\alpha = 0.252$) at $Y = 7.1$; this is the most uncertain (or “most fractal”) part of the shadow. Beyond this value, there is a discontinuity in the uncertainty exponent; we see that $\alpha = 1$ for $Y > 7.1$, which is the end of the black hole shadow. The end of the shadow corresponds to $p_\phi = p_\phi^*$: a ray incident from $r \rightarrow \infty$ with $p_\phi \geq p_\phi^*$ is not permitted to fall into the black holes by the critical contour $h(\rho, z) = p_\phi$. The value of Y corresponding to p_ϕ^* is found to be $Y = 7.13$. This agrees well with the results of Figures 5.5(c) and 5.5(d).

In the case $d = 1$, the one-dimensional shadows become more uncertain as one increases the value of Y (for $Y < 7.13$). This may be due to the existence of a “pocket” with three

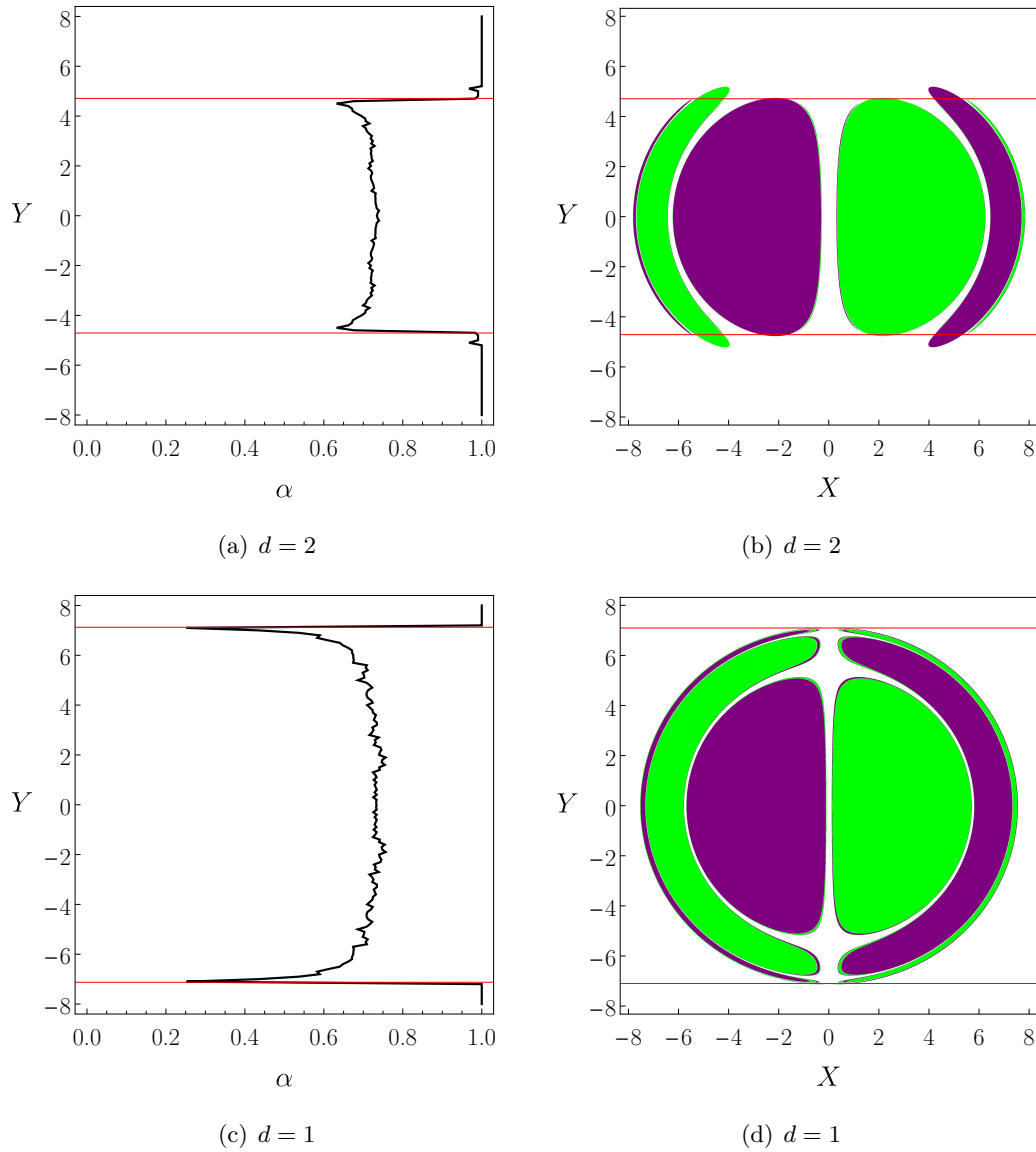


Figure 5.5. Uncertainty exponent $\alpha(Y)$ for one-dimensional slices ($Y = \text{constant}$) of Majumdar–Papapetrou di-hole shadows. (a)–(b) For $d = 2$, the uncertainty exponent is able to distinguish between the regular and fractal (Cantor-like) parts of the shadow. The critical value of Y , at which the type II/III fundamental orbits coincide, is shown as a red horizontal line. (c)–(d) In the case $d = 1$, the uncertainty exponent confirms that the entire shadow boundary is fractal. The end of the shadow, corresponding to $p_\phi = p_\phi^*$, occurs at $Y = 7.13$. This is shown as a red horizontal line. We have transposed the graphs of α in (a) and (c) to aid comparison with the shadow images in (b) and (d).

escapes, leading to the black holes and to spatial infinity; see the discussion in Section 4.7.1. The escapes become narrower as one increases p_ϕ towards p_ϕ^* . This results in greater final-state unpredictability, because the “pocket” acts as a randomising region.

For the highly symmetric case $d = 1$, as $p_\phi \rightarrow p_\phi^*$, the width of the three escapes in phase space tends to zero. In this limit, we also anticipate that $\alpha \rightarrow 0$. In principle, this could be the case; however, the computation of α is difficult in this regime, due to the fact that the volume of the shadow tends to zero. This is similar to the limit $\gamma \rightarrow 0$ for the γ -Cantor basins, described in Section 5.4.1.

5.5 The Wada property

5.5.1 Wada basins

A collection of three or more open sets $S = \{B_i | i \in \{1, 2, \dots, N\}, N \geq 3\}$ are said to satisfy the *Wada property* if each point p on the boundary of any of the $B_i \in S$ is on the boundary of *all* $B_i \in S$. Such sets are difficult to visualise, and at first seem impossible to realise; however, the first explicit example of sets which possess this property – the *lakes of Wada* – was given by Yoneyama in 1917 [244].

The lakes of Wada are constructed using the following simple geometrical procedure [244, 253]. We start with a double annulus; this is taken to be an island in a red ocean with two lakes, one of which is filled with green water, the other with blue. At time $t = 0$, we dig a canal from the red ocean, such that each point on the land is within a distance of 1 unit of red ocean water. At time $t = \frac{1}{2}$, we dig a canal from the green lake, which brings green water to within a distance $\frac{1}{2}$ from each point on the land. At time $t = \frac{3}{4}$, we dig a canal from the blue lake so that each point on the land is within a distance of $\frac{1}{3}$ from blue water. At time $\frac{7}{8}$, a canal is dug from the first canal so that each point on land is within a distance $\frac{1}{4}$ of every point on land. We continue in this fashion until time $t = 1$, when the remaining land is a continuum in the plane of zero measure which bounds *three* open sets, i.e., the red ocean, the green lake and the blue lake. The boundary of the lakes of Wada is an example of an *indecomposable continuum*, a concept first described by Brouwer [254].

The Wada property began as a topological curiosity; however, Kennedy and Yorke [245] demonstrated that is quite common in dynamics, arising generically in the basins of systems with three or more attractors or escapes. Recall from Section 2.6.5 that a point p is a *boundary point* of the basin B if every open neighbourhood of p contains at least one point in B and one point not in B . If a point p is a boundary point of $N \geq 3$ basins B_i , $i \in \{1, 2, \dots, N\}$, then p is said to be a *Wada point*. If all the boundary points of a basin B are Wada points, then B is called a *Wada basin*. The Wada property may be rephrased as follows: a point p on the basin boundary is a Wada point if and only if every open neighbourhood of p has a non-empty intersection with at least three different basins (open sets).

In their seminal paper, Kennedy and Yorke [245] argue that the Poincaré return map of the forced damped pendulum with four attractors exhibits the Wada property. The

Wada property has since been revealed, using a range of techniques (see Section 5.6), in a variety of dynamical systems, such as the forced damped pendulum, the three-disc model, the Hénon–Heiles Hamiltonian system, and others [112].

Here, we are primarily concerned with open Hamiltonian systems, so the subsequent discussion will refer only to exit basins; however, this could be generalised to include basins of attraction for dissipative systems. When multiple escapes are present in phase space, we typically wish to determine to which basin a point belongs, i.e., through which escape the particle will leave the system when integrated forwards in time. If the basin boundary is sufficiently complicated, a small uncertainty in fixing the initial conditions may result in a large uncertainty in the final state of the system. In situations where the basin boundary is a Wada boundary, one encounters a high level of indeterminacy and an extreme sensitive dependence on initial conditions, despite the underlying system being fully deterministic. The presence of Wada basins in phase space is therefore a hallmark of deterministic chaos in non-linear dynamical systems.

5.5.2 The Wada property in Cantor basins

It is clear from Yoneyama’s geometrical construction of the lakes of Wada, reviewed in Section 5.5, that the Wada property is intimately related to iteration. Furthermore, we have seen that a Cantor set can be generated using a similar iterative procedure to that which is used to construct the lakes of Wada. In this section, we demonstrate how the Wada property can arise generically in basins with a Cantor-like boundary.

Recall the Cantor basins presented in Section 4.6.2. One can construct a set of $N \geq 3$ basins $\{B_i\}_{i=1}^N$ from the complement of the middle-thirds Cantor set \mathcal{C} in $I = [0, 1]$ as follows. Denote by J_k the union of the 2^{k-1} open intervals which are *removed* from I at the k th iteration in the construction of \mathcal{C} . Now, add J_k to the basin B_i , where $k = i \bmod N$. The N basins can then be expressed as an infinite union of disjoint open sets:

$$B_i = \bigcup_{k=i \bmod N} J_k, \quad i \in \{1, 2, \dots, N\}. \quad (5.9)$$

Let us colour the open intervals which are removed at each step using N colours, i.e., we assign the intervals a label from the set $\{c_1, c_2, \dots, c_N\}$. The label is chosen in accordance with the iteration: all intervals which belong to the i th basin are labelled by the colour c_i . Clearly, $\partial B_1 = \partial B_2 = \dots = \partial B_N = \mathcal{C}$; i.e., all of the basins B_i share a common boundary, namely the Cantor set. Our aim is to demonstrate that the basins $\{B_i\}_{i=1}^N$ possess the Wada property: every open neighbourhood of a point which belongs to the basin boundary has a non-empty intersection with all N different basins.

As described in Section 4.6.2, the construction of the Cantor basins can be understood in terms of symbolic dynamics for the Cantor set, where each interval is encoded using a string of symbols from a symbolic alphabet $\mathcal{A} = \{0, 1, 2\}$. At each iteration, the left-hand (closed) interval is assigned the digit 0; the middle (open) interval is assigned the digit 1; and the right-hand (closed) interval is assigned the digit 2. The intervals which are assigned a colour (i.e., which belong to an exit basin) are all those whose symbolic representation ends in the digit 1. We iterate on the remaining intervals (corresponding to the digits 0 and 2) *ad infinitum*.

The boundary points are represented in symbolic dynamics by sequences of infinite length which do not contain the symbol 1. We note that sequences which terminate in the digit 1 correspond to *open* intervals, i.e., they do not contain the endpoints. Let $S_{(k)}$ be a sequence of length k which does not contain the symbol 1. The endpoints of the interval with representation $S_{(k)}1$ have symbolic representation $S_{(k)}1\bar{0}$ and $S_{(k)}1\bar{2}$. However, these points also admit alternative representations in symbolic dynamics, namely $S_k0\bar{2}$ and $S_k2\bar{0}$, respectively. Clearly, the latter representation shows that these points are boundary points, and should not be included in the interval $S_{(k)}1$. Thus $S_{(k)}1$ is open.

The open intervals in the basin B_i (which are labelled using the colour c_i) are represented in symbolic dynamics by sequences of length l , with $l = i \bmod N$, which terminate in the symbol 1, and whose first $l - 1$ digits do not contain the symbol 1. For example, the open interval labelled 020021 (which can be read as left-right-left-left-right-middle) belongs to the basin B_i , where $6 = i \bmod N$.

It is possible to use the symbolic dynamics for the Cantor set to demonstrate that the N -colour Cantor basins possess the Wada property. Given any boundary point $p \in \mathcal{C}$, whose symbolic representation is of the form $p_1p_2p_3\dots$, where each $p_i \in \{0, 2\}$, one can always find a point in *any* of the N basins which is arbitrarily close to p . For if we are given a point $q \in B_i$, with symbolic representation $p_1p_2\dots p_l1$, it is always possible to find a point $q' \in B_i$ with symbolic representation $p_1p_2\dots p_{l'}1$, where $l' = l \bmod N$ and $l' > l$. The point q' is closer to p than q is, but is still in the basin B_i . Since p and B_i are arbitrary, this demonstrates that every open neighbourhood of each boundary point p has non-empty intersection with *all* of the N basins $\{B_1, B_2, \dots, B_N\}$. Thus, the N -colour Cantor basins (with $N \geq 3$) possess the Wada property. This simple argument illustrates how the Wada property may arise generically in fractal sets with Cantor-like structure, which are typically constructed using iterative processes.

5.6 Numerical test for Wada basins

5.6.1 The merging method

As described in Section 5.5, the Wada property arises naturally in non-linear dynamical systems. A number of methods have been proposed to test this striking property. We present a brief overview of these methods here, before reviewing the *merging method* – a recently developed algorithm which tests whether an exit basin diagram exhibits the Wada property.

In 1996, Nusse and Yorke [255] established that an unstable manifold which crosses three (or more) basins in phase space could be used to prove the existence of Wada boundaries. The main goal of the work of Nusse and Yorke was to develop criteria for Wada basins which could be verified numerically; Nusse and Yorke applied their method to verify the Wada property in the basins of the Hénon map and the forced damped pendulum [255]. The so-called *Nusse–Yorke method* has since been used to demonstrate that the basins of Gaspard–Rice three-disc model possess the Wada property [226]. The Nusse–Yorke method involves the computation of the unstable manifold of a saddle point

which lies on the basin boundary. The algorithm therefore requires a detailed knowledge of the dynamical system under consideration, as well as the computation of unstable manifolds; this is often technically and computationally demanding. Indeed, many studies [226, 246, 247, 256] have been devoted to applying the Nusse–Yorke method to a given dynamical system for a particular set of parameters.

In 2015, Daza *et al.* [257] proposed a new computational method to verify the existence of Wada basins in dynamical systems. This approach, referred to as the *grid method*, is based on successive refinements of a grid in phase space. At each stage of the algorithm, further trajectories are computed at higher and higher resolution. This method can be automated in a simple fashion, and does not require a detailed knowledge of the underlying dynamical system. The grid method is therefore easily applied to any dynamical system, and can test for the presence of Wada basins up to finite resolution. One drawback of this approach is that it requires the computation of new trajectories of the system at high resolution. It is therefore computationally expensive.

More recently, with the principal aim of overcoming the drawbacks of the Nusse–Yorke method and the grid method, Daza *et al.* [248] developed a novel method to test for Wada basins up to finite resolution. The method, referred to as the *merging method*, does not require a detailed knowledge of the underlying dynamical system, and is computationally inexpensive. Of the three methods outlined here, the merging method is the fastest, and the only one able to provide a reliable test of the existence of Wada basins through a straightforward examination of the basins (at finite resolution), without the need to compute new trajectories of the system or its invariant manifolds.

The method is based on a simple but counterintuitive observation: Wada basin boundaries are invariant under the action of merging any two of the basins together. To see this, consider $N \geq 3$ disjoint open sets (i.e., basins) $\{B_1, B_2, \dots, B_N\}$. Suppose $p \in \partial B_i$ for all $i \in \{1, 2, \dots, N\}$; in other words, p is a Wada point. Define a new basin $B_{1,2} = B_1 \cup B_2$ (that is, *merge* the basins B_1 and B_2 together). Clearly $p \in \partial B_{1,2}$ and $p \in \partial B_j$ for all $j \in \{3, \dots, N\}$, so p is still a boundary point of the merged basins. Since the choice of basins B_1 and B_2 was arbitrary, the action of merging any two of the basins together leaves the Wada boundary points invariant. In fact, one could continue merging the basins together in this fashion, until one is left with a pair of basins B_i and $\bigcup_{j \neq i} B_j$ for some $i \in \{1, 2, \dots, N\}$. Clearly $p \in \partial B_i$ and $p \in \partial \left(\bigcup_{j \neq i} B_j \right)$. Figure 5.6 illustrates this striking property. In Figures 5.6(a)–5.6(d), we show the exit basins in phase space and their merged versions for the Majumdar–Papapetrou di-hole with initial conditions taken in the (ρ, z) -plane and an angular momentum parameter $\Delta p_\phi = p_\phi^* - p_\phi = 0.05$ (see Section 5.2.2). On first inspection, it appears that we have simply changed the colours used to depict the basins. However, closer examination reveals that the boundaries in each figure are the same, but that two of the three basins have been merged to form a new basin in each case. It is straightforward to see that non-Wada boundaries do not exhibit this invariance under the pairwise merging of the basins. Figures 5.6(e)–5.6(h) depict a simple example of regular (i.e., non-fractal) basins – constructed by partitioning the unit disc into three equal sectors – and their merged counterparts. For this example, only one of the boundary points (the centre of the disc) is a Wada point. The non-Wada boundaries

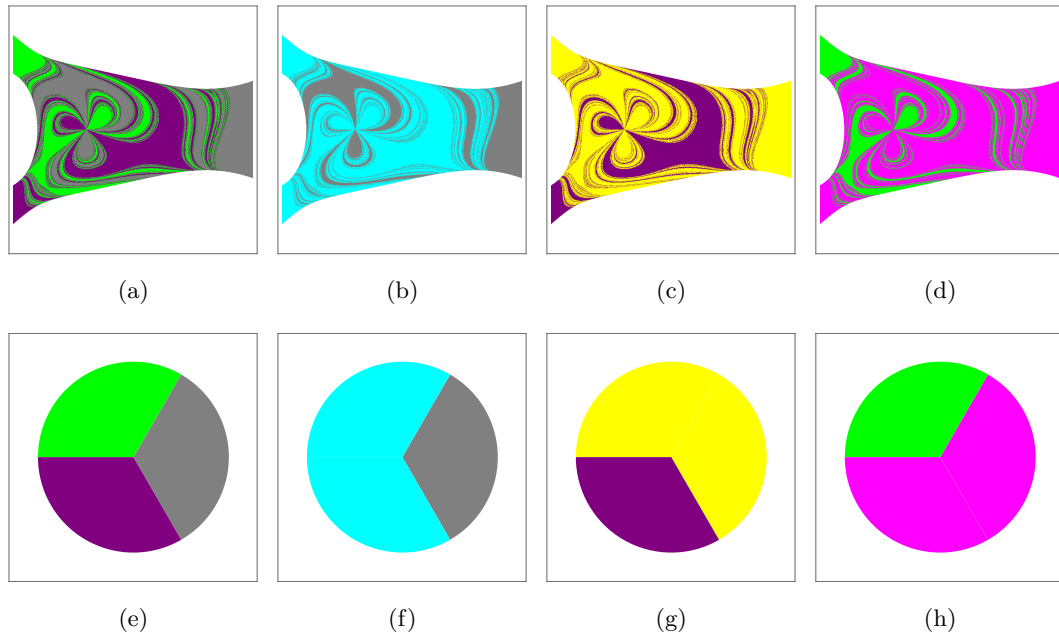


Figure 5.6. Illustration of the merging method for Wada basins and regular basins. (a)–(d) Wada basins for the Majumdar–Papapetrou di-hole system for initial conditions in the (ρ, z) -plane with $\Delta p_\phi = 0.05$. (a) The original three-colour basins: rays which plunge into the upper (lower) black hole are coloured green (purple); rays which escape to infinity are coloured grey. (b) The basins corresponding to the black holes are merged [cyan]. (c) The basins corresponding to the upper black hole and spatial infinity are merged [yellow]. (d) The basins corresponding to the lower black hole and spatial infinity are merged [magenta]. (e)–(h) A simple example of regular basins and their merged versions. In this case, only the centre of the disc, where the three basins meet, is invariant under the pairwise merging of the basins; this is a Wada point. The remaining boundary points are not Wada points.

are destroyed by the pairwise merging of the basins.

Mathematically, to check that the basins of a system satisfy the Wada property, it suffices to check that the boundary remains unaltered under the merging of the basins. However, this procedure would require either (i) a closed-form expression for the basins and their boundaries; or (ii) an image of the basins up to arbitrarily high resolution.

Typically, the exit basins of a dynamical system are computed using a finite grid of initial conditions in the (u, v) -plane, where u and v are the coordinates of a two-dimensional surface of section in phase space. The basins can, of course, be realised in higher dimensions; however, we restrict our attention to the two-dimensional case for simplicity. Here, a “pixel” is taken to be a square of side length L centred on the initial condition (u_0, v_0) . The exit basins are realised by colouring the pixels according to the final state of the trajectory with initial condition given by the point (u_0, v_0) , i.e., the centre of the pixel.

Due to the issues caused by the finite resolution of the grid used to realise the basins, we *fatten* the basin boundaries by a factor of the *fattening parameter* r ; that is, we replace each pixel belonging to the (original) *slim boundaries* by itself plus its r nearest neighbours.

For a square grid of initial data, this is typically achieved using the Chebyshev metric; however, other grid layouts and metrics can be employed without drastically affecting the method.

Having computed the boundaries and applied the fattening procedure for some fattening parameter r , we are left with the set of slim boundaries $\{\partial B_i\}_{i=1}^N$ and their fattened versions $\{\partial_r B_i\}_{i=1}^N$. The condition of the method states that if all of the (original) slim boundaries are contained in all of the fat boundaries, i.e., $\partial B_i \subset \partial_r B_j$ for all $i, j \in \{1, 2, \dots, N\}$, then the basin boundary is a Wada boundary up to the resolution of the fattening parameter r . The method is reliant on the choice of parameter r ; we begin with $r = 1$ and increase its value until either all of the basin boundary points are classified as Wada points, or a halting condition $r > r_{\text{stop}}$ is met. Of course, the merging method only ascertains whether or not a basin is Wada up to a resolution determined by the internal parameters of the method: the grid size L , and the fattening parameter r .

Let us summarise the steps of the merging method for clarity.

1. Begin with a finite resolution image of N exit basins $\{B_i\}_{i=1}^N$. Recall that the method does not require any prior knowledge of the dynamical system; however, this will be necessary to compute the basins in the first place.
2. Choose one basin B_i for some $i \in \{1, 2, \dots, N\}$, and merge all of the other basins $\bigcup_{j \neq i} B_j$. This yields a two-colour exit basin diagram. Repeat this procedure for each basin B_i in turn. This results in a set of N two-colour basin diagrams (see Figure 5.6).
3. Determine the boundary of each two-colour basin diagram. Practically, this is achieved by identifying pixels which have at least one neighbour of the opposite colour. This yields a set of N slim boundaries $\{\partial B_i\}_{i=1}^N$.
4. Starting with $r = 1$, fatten each slim boundary by a factor of r (i.e., take each boundary point and all other points within a distance r) to obtain N fat boundaries $\{\partial_r B_i\}_{i=1}^N$.
5. Consider a fat boundary $\partial_r B_j$ for some $j \in \{1, 2, \dots, N\}$. Check whether the union of slim boundaries is contained inside the fat boundary. Repeat for each of the N fat boundaries $\partial_r B_j$ in turn. If the union of the slim boundaries is contained inside *every* fat boundary, then the basins are said to possess the Wada property, up to the resolution determined by the fattening parameter r . If this is not the case, increase the value of r ($r \mapsto r + 1$) and return to step 4 of the algorithm. Repeat this process until the halting condition $r > r_{\text{stop}}$ is satisfied. If this condition is met and the union of slim boundaries is not contained in each fat boundary, then the merging method classifies the basins as non-Wada. In the case of partially Wada boundaries (i.e., some boundary points are Wada points, but others are not), the merging method provides a list of all of the boundary points in the original image and their classification as Wada or non-Wada.

Daza *et al.* [248] present an overview of the merging method, and provide a detailed analysis of the method. Moreover, the method is used to test for the Wada property in

the basins of two well-known systems: the Hénon–Heiles Hamiltonian; and the Newton method to find the n th roots of unity in the complex plane. Numerical investigation of these systems demonstrates that the merging method correctly classifies basins which are known to possess the Wada property; and that the classification using a fattening parameter r is independent of (i) the fractal dimension of the boundary, and (ii) the number of basins.

5.6.2 The Wada property in Majumdar–Papapetrou exit basins

As described in Sections 5.2.2 and 5.6.1, the exit basins in the phase space of the Hénon–Heiles Hamiltonian system [117] are known to possess the Wada property, as demonstrated by Aguirre *et al.* [246], who employed the Nusse–Yorke method [255]. This has since been confirmed using the grid method [257] and the merging method [249]. Given the striking similarities between the Hénon–Heiles system and the Majumdar–Papapetrou di-hole system (see Section 5.2.2 and Appendix C), one would naturally expect the basins of the Majumdar–Papapetrou di-hole – shown in Figure 5.2 – to exhibit the Wada property. Here, we test this claim using the merging method.

We applied the merging algorithm described in Section 5.6.1 to the exit basins of the Majumdar–Papapetrou system with initial data taken in both the (ρ, z) -plane and the (ρ, p_ρ) -plane (see Figures 5.2(a)–5.2(c) and Figures 5.2(d)–5.2(f), respectively), where the basins were realised using a grid of $10^3 \times 10^3$ equally spaced initial conditions. We tested all boundary points for the Wada property for a selection of exit basin diagrams with $\Delta p_\phi \in [0.02, 0.15]$, where $\Delta p_\phi = p_\phi^* - p_\phi$. For this choice of parameters, the merging algorithm classified *all* boundary points as Wada points for a fattening parameter of $r = 3$.

We remark that, as $\Delta p_\phi \rightarrow 0$ (i.e., as $p_\phi \rightarrow p_\phi^*$), the width of the three escapes in phase space tends to zero. Moreover, KAM islands become dominant – as shown in Figures 5.2(c) and 5.2(f) – and the escape time for photons which start inside the scattering region blows up [227]. It is therefore computationally expensive to verify the Wada property for small values of the parameter Δp_ϕ . Although we have not directly verified the Wada property for Majumdar–Papapetrou basins with $\Delta p_\phi < 0.02$, we expect all boundary points to retain the Wada property in the limit of small escapes ($\Delta p_\phi \rightarrow 0$). A detailed discussion of the limit of small escapes in open Hamiltonian systems is given by Aguirre and Sanjuán [227].

5.6.3 The Wada property in Majumdar–Papapetrou shadows

We now turn our attention to the shadows of the Majumdar–Papapetrou di-hole, described in detail in Section 4.7.2, and reviewed in this chapter in Section 5.3; see Figure 4.17 for examples of shadows of equal-mass ($M_\pm = M = 1$) black holes separated by distances $d = 1$ and $d = 2$.

We applied the merging method (Section 5.6.1) to test for the Wada property in Majumdar–Papapetrou di-hole shadows for a selection of coordinate separations in the range $d \in [0, 3]$. In each case, we employed a ray-tracing method (Section 4.7.2) to generate black hole shadow images for an observer situated at a radial distance $r_0 = 50M$ with

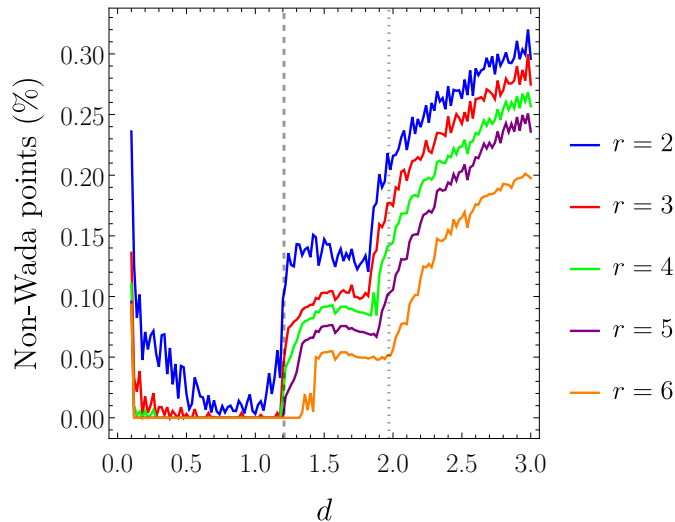


Figure 5.7. Results of the merging method for the equal-mass Majumdar–Papapetrou di-hole with coordinate separation d . The numerical algorithm was performed for a range of values of the fattening parameter r . The dashed vertical line indicates the critical separation $d = \hat{d} = 1.2085$, below which the shadows exhibit the full Wada property. The dotted vertical line at $d \approx 1.97$ marks a second transition in the nature of the shadow, where there is an increase in the number of non-Wada points.

a viewing angle of $\theta = \frac{\pi}{2}$. Practically, this involved numerically integrating Hamilton’s equations for the Hamiltonian (5.2) for a $10^3 \times 10^3$ grid of evenly spaced initial conditions in the (X, Y) -plane (the observer’s image plane).

The results of the merging algorithm are shown in Figure 5.7, which depicts the percentage of boundary points which are *not* classified as Wada points by the algorithm, as the coordinate separation d is varied. The results of the algorithm provide evidence that the shadows are fully Wada (i.e., all boundary points are Wada points) for separations in the range $0.1 \lesssim d \lesssim 1.2$. Moreover, the results of Figure 5.7 suggest that there is a qualitative transition in the nature of the shadows at $d \approx 1.2$, after which the shadow is only partially Wada.

Such a qualitative change was anticipated in our study of the fundamental orbits in Sections 5.3.2 and 5.3.3. We argued that, for a fixed value of p_ϕ , the existence of three types of photon orbits, as shown in Figure 4.11, gives rise to Cantor-like structure in the corresponding one-dimensional shadows (see Chapter 4). In Section 5.5.2, we presented a simple argument based on symbolic dynamics which illustrates that a set of $N \geq 3$ exit basins which have a Cantor-like boundary will possess the Wada property. Hence, if all three types of fundamental orbits exist for $0 \leq p_\phi < p_\phi^*$, then we anticipate that the two-dimensional shadow will be totally Wada. If, on the other hand, there exists some value $\hat{p}_\phi < p_\phi^*$ for which the type II and III fundamental orbits cease to exist – as shown in Figure 4.11(c) – then the one-dimensional shadows with $\hat{p}_\phi < p_\phi < p_\phi^*$ will *not* be fractal. As a result, the two-dimensional shadow will only be partially Wada: the boundary points corresponding to $0 \leq p_\phi \leq \hat{p}_\phi$ will be Wada points, whereas those corresponding to $\hat{p}_\phi < p_\phi < p_\phi^*$ will not. In Section 5.3.3, we demonstrated that the latter

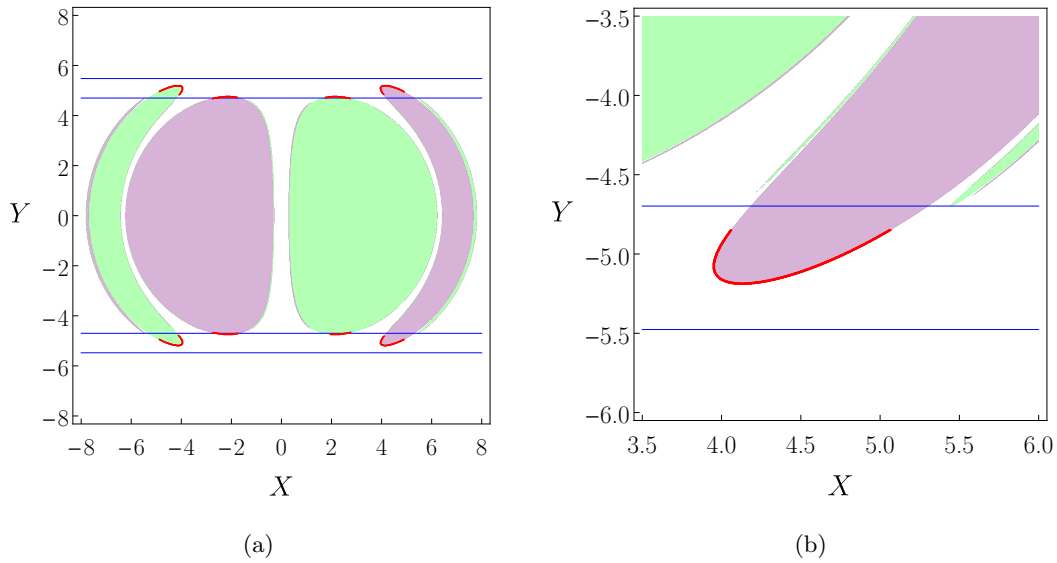


Figure 5.8. Results of the merging method to test for Wada boundary points in the equal-mass Majumdar–Papapetrou di-hole shadow with coordinate separation $d = 2M$. (a) Shadow of the Majumdar–Papapetrou di-hole in the image plane with coordinates (X, Y) . The shadow of the upper (lower) black hole is shown in green (purple). The non-Wada boundary points found by the merging method (with $r = 5$) are shown in red; critical values of Y , which delimit the non-fractal regions, are shown in blue. The latter are obtained by looking at critical values of p_ϕ . (b) Magnified region of the shadow (a) which shows good agreement between the non-Wada points detected by the merging method and the non-fractal regions of the shadow.

scenario occurs for sufficiently separated black holes with $d > \hat{d} \approx 1.2085$. This prediction matches the observed transition in the results of the merging algorithm, shown in Figure 5.7.

Figure 5.7 also indicates that there is a second qualitative change in the nature of the shadow structure at $d \approx 1.9$. It is likely that this transition occurs when the top of the main globular shadow components impinge upon the regular region of the shadow. For $d \gtrsim \hat{d}$, only the tips of the primary eyebrow-like components are regular. As one increases d , the regular region of the shadow (i.e., the domain in Y for which all one-dimensional shadows are non-fractal) begins to incorporate the top of the globular features (main shadows). Numerical investigations of the shadow structure and the location of the regular region suggest that this occurs at $d \approx 1.97$, which agrees well with the observed transition in Figure 5.7 (shown as a dotted vertical line).

To confirm our interpretation of the Wada property in Majumdar–Papapetrou di-hole shadows, we illustrate the results of the merging method in Figure 5.8. We consider a Majumdar–Papapetrou di-hole shadows seen by an observer with a viewing angle of $\theta = \frac{\pi}{2}$, for black holes of mass M separated by coordinate distance $d = 2M$. The exit basin corresponding to the upper (lower) black hole are plotted in green (purple); the exit basin corresponding to spatial infinity is plotted in white. The non-Wada points detected by the merging method (for a fattening parameter $r = 5$) are shown in red. The regular (non-fractal) region of the shadow is delimited by the blue horizontal lines; this

region is determined by considering the values of Y corresponding to $p_\phi \in (\hat{p}_\phi, p_\phi^*)$. The results of Figure 5.8 confirm that all of the non-Wada points lie within the regular region. Furthermore, one can see from Figure 5.8(a) that the regular zone has begun to impinge on the primary shadows for $d = 2 > 1.97$.

The agreement between the boundaries of the regular region [blue curves] and the non-Wada points [red] detected by the merging method could be improved by (i) increasing the resolution of the exit basin diagram; and (ii) taking $r_{\max} \rightarrow \infty$, where r_{\max} is the location of the observer. We caution that implementing these proposals would make the merging method more computationally expensive.

5.7 Discussion

5.7.1 Extensions

Do the shadows cast by real binary black holes, such as the progenitors of gravitational waves detected by LIGO–Virgo, exhibit the Wada property? This remains an open question. Despite the fact that singleton black hole systems have been successfully imaged [1], there appears to be little prospect of imaging a binary black hole system in the near future. However, realistic simulations of gravitational lensing by dynamical binary black holes have been conducted by Bohn *et al.* [190]. Classifying these high-resolution shadow images as fully Wada, partially Wada or otherwise by applying the merging method would certainly be of theoretical interest. One could also apply the algorithm to shadows in other binary systems, such as the imitative models explored in the recent work of Cunha *et al.* [234].

Aguirre and Sanjuán [227] explore the exit basins of open Hamiltonian systems in the limit that the escape width tends to zero, using the Gaspard–Rice three-disc model and the Hénon–Heiles Hamiltonian system as illustrative cases. They find that, as the width of the escapes decreases, the basins become *totally uncertain* (i.e., the uncertainty exponent approaches zero with the width of the escapes). These totally uncertain basins resemble the case of *riddled basins* in dissipative systems. A basin is said to be *riddled* if all of its points are arbitrarily close to points from another basin [258]. Riddled basins coincide with their boundary.

In the case of the Majumdar–Papapetrou Hamiltonian system with $d = M$ (highly symmetric case), we have three exits in phase space whose widths approach zero in the limit $\Delta p_\phi \rightarrow 0$. We anticipate that the exit basins in phase space and the one-dimensional shadows would become totally uncertain in the limit $\Delta p_\phi \rightarrow 0$. The basins of Figure 5.2, the black hole shadows of Figure 5.5(c) and the uncertainty exponent in Figure 5.5(d) all support this claim; however, an accurate calculation of the uncertainty exponent in the limit $\Delta p_\phi \rightarrow 0$ would be computationally expensive.

5.7.2 Conclusions

In this chapter, we have applied techniques from the field of non-linear dynamics to study fractal structures which arise in the shadows of a static binary black hole model – the Majumdar–Papapetrou di-hole solution – which is an exact solution to the Einstein–

Maxwell equations of gravity and electromagnetism. The solution can be regarded as a surrogate for a more “realistic” dynamical binary black hole system, such the progenitors of the gravitational waves detected by the LIGO–Virgo collaboration.

The black hole shadow is intimately related to the scattering of photons, which follow null geodesics on the curved spacetime geometry. The null geodesics of the Majumdar–Papapetrou geometry are governed by a two-dimensional Hamiltonian dynamical system with a conserved parameter for rays (the photon’s azimuthal angular momentum). Intriguingly, this Hamiltonian system shares many qualitative features with the well-known Hénon–Heiles system [117], a paradigmatic two-dimensional time-independent model for chaotic scattering in Hamiltonian dynamics.

In Chapter 4, we saw that the Majumdar–Papapetrou di-hole system is an example of chaotic scattering: the existence of three dynamically connected fundamental photon orbits gives rise to a fractal set of scattering singularities of measure zero. The chaotic scattering of photons by the two black holes is responsible for Cantor-like fractal structure in the binary black hole shadows. Our perspective is that black hole shadows can be viewed as exit basins defined with respect to the image plane of a distant observer. Black hole shadows are therefore amenable to techniques from non-linear dynamics which have been developed to characterise and quantify fractal structures.

In Section 5.2.1, we reviewed the key features of the Majumdar–Papapetrou di-hole geometry and the Hamiltonian formalism for null geodesics (first covered in Chapter 4). Taking inspiration from the study of the Hénon–Heiles system in [246], we analysed the geodesic dynamics of the Majumdar–Papapetrou di-hole by considering exit basins in phase space in Section 5.2.2. Then, in Section 5.3, we reviewed the black hole shadows of the equal-mass Majumdar–Papapetrou di-hole system for different values of the coordinate separation d between the black holes. The role of fundamental photon orbits on the variation in shadow structure was discussed (Section 5.3.2). We outlined a numerical method which could be used to calculate the critical value of the separation parameter \hat{d} (Section 5.3.3); we conjectured that Majumdar–Papapetrou di-hole shadows with $d < \hat{d}$ will be fully Wada, whereas shadows with $d > \hat{d}$ will only be partially Wada (i.e., they will exhibit both Wada and non-Wada regions). In Section 5.4.1, we reviewed the notion of Hausdorff dimension for fractal sets and the uncertainty exponent, first introduced by Grebogi *et al.* [251] as a quantitative measure for the indeterminacy in final-state prediction when fractal boundaries are present in phase space. We presented a numerical method which can be employed to compute the uncertainty exponent for exit basins. To test and calibrate the numerical method, we calculated the uncertainty exponent for the Cantor basins (first introduced in Section 4.6.2) and compared the numerical data to known analytical results (Section 5.4.2). We then applied the numerical method to one-dimensional Majumdar–Papapetrou di-hole shadows (Section 5.4.3); the results support our understanding of the shadow structure and fundamental photon orbits, outlined in Chapter 4 and Section 5.3. In Section 5.5.1, we reviewed the Wada property, with an emphasis on the exit basins of open Hamiltonian systems. Using a simple argument based on “symbolic dynamics” for Cantor basins, we demonstrate that the Wada property arises generically in systems with three or more exit basins who share a common Cantor-like

fractal basin boundary (Section 5.5.2). We applied the merging method (reviewed in Section 5.6.1) to test for the Wada property in (i) the exit basins in phase space of the Majumdar–Papapetrou di-hole (Section 5.6.2); and (ii) the shadows cast by the black holes, i.e., exit basins on the image plane of an observer (Section 5.6.3).

Using the merging method, we have shown that the Majumdar–Papapetrou di-hole shadows can exhibit either the partial Wada property or the full Wada property, depending on the separation of centres d . In the exit basins of open Hamiltonian systems, the Wada property is associated with high levels of indeterminacy and unpredictability, despite the underlying dynamics being deterministic. In the case of a pair of black holes, the existence of Wada basins implies that a photon whose initial conditions are chosen close to a Wada boundary can end up in one of three final states: it can plunge into either of the black holes or escape to spatial infinity.

An important property of the merging method of [248] is that it is agnostic to the underlying dynamical system: no knowledge of the system or its invariant sets is required; all that is necessary is an exit basin diagram at finite resolution. A key result of this work is that the algorithm successfully detected a “phase transition” in the black hole shadow, as it changed from fully Wada to partially Wada at a certain value of the separation parameter d . The numerical results of the merging method agree well with our theoretical predictions based on fundamental photon orbits (Sections 5.3.2 and 5.3.3). In the exit basins of dynamical systems where the underlying dynamics is either unknown or too complicated to study analytically, the merging method can provide novel physical insight.

As described in Section 5.6.1, a number of computational methods have been proposed to test for Wada basins, including the Nusse–Yorke method [255]. The merging method, employed in this work, offers several advantages over the Nusse–Yorke method. In the latter approach, to verify the Wada property, one must compute an unstable manifold and verify that it crosses all of the exit basins in phase space. In the case of black hole shadows for the Majumdar–Papapetrou di-hole, the image plane mixes phase space and parameter space: the image plane coordinates are dependent on the independent phase space coordinates and the conserved parameter p_ϕ . It is not clear how one would construct an unstable manifold on the image plane; as a result, it does not appear possible that the Nusse–Yorke method could be used to test for the Wada property in Majumdar–Papapetrou di-hole shadows.

The work presented in this chapter (based on that of [249]) is, to our knowledge, the first demonstration of the Wada property in the exit basins and/or the black hole shadows of a general relativistic system. Our work demonstrates that tools from the field of non-linear dynamics (e.g. the uncertainty exponent and the merging method to test for Wada basins) can be exploited to deepen our understanding of scattering processes in general relativity. Furthermore, this investigation of the Majumdar–Papapetrou di-hole system shows that there exist novel dynamical systems in general relativity (and in the wider field of gravitational physics) which are amenable to techniques from non-linear dynamics and chaos theory.

Chapter 6

Stable photon orbits in stationary axisymmetric spacetimes

6.1 Introduction

The epoch-making gravitational-wave detections by the LIGO–Virgo collaboration provide compelling evidence for the existence of stellar-mass black holes in nature [18]. Moreover, the characteristic “chirp” profiles of these events are consistent with the inspiral, merger and ringdown phases of a binary black hole coalescence, which settles down to an equilibrium state described by Kerr’s solution of general relativity [35].

The ringdown phase of a binary black hole coalescence provides strong evidence that the end-product of the merger possesses a *light-ring* [259], i.e., a family of *unstable* photon orbits [76,177]. The parameters of the light-ring – in particular, its orbital frequencies and Lyapunov exponents – are related to the frequency and decay rate of the ringdown phase of the gravitational-wave signal [260–264]. If the merged body is endowed with a light-ring, this will be associated with a range of observational phenomena, including multiple lensing images of distant sources [7]; diffraction effects such as glories and orbiting [265]; and a characteristic spectrum of quasinormal modes [266].

Intriguingly, the correspondence between the object’s complex quasinormal mode oscillations and the existence of an unstable photon orbit suggests that any object with a light-ring will initially vibrate like a black hole [259]. As a result, the LIGO–Virgo gravitational-wave detections do not rule out the possibility that the merged body possesses a light-ring but *not* an event horizon. Objects with a light-ring are dubbed *ultra-compact* in the literature [267,268].

Horizonless ultracompact objects – such as boson stars [269], Proca stars [270], gravastars [271,272], and wormholes [273] – possess unstable light-rings which are generically accompanied by inner *stable* photon orbits [267]; the latter are associated with a range of phenomenological features. First, the existence of stable photon orbits allows for the trapping and storage of energy in bounded regions; this allows instabilities (e.g. fragmentation and collapse [267] or ergoregion instabilities [274,275]) to flourish. Second, Keir [276] found that linear waves – which serve as a model for non-linear perturbations – decay (no faster than) logarithmically with time, which dominates power-law Price decay [277];

this slow decay suggests that ultra-compact objects are unstable at the non-linear level. Finally, ultra-compact objects with a stable photon orbit exhibit a modified late-time gravitational-wave ringdown which depends on the internal structure of the compact object, and differs from that of a black hole [259]. As we enter a new era of precision gravitational-wave experiments, this modified ringdown will be tested and constrained.

Various strands of work suggest that stable photon orbits are relevant not only in horizonless geometries, but also play a key role in black hole spacetimes. In four dimensions, stable photon orbits have been shown to exist within the inner horizons of Kerr–Newman black holes and around naked singularities [77, 78, 278–283]; and around black holes (or solitons) with a non-zero cosmological constant [284, 285]. In higher-dimensional contexts, stable photon orbits have been revealed in the exterior regions (i.e., outside the event horizon) of five-dimensional black rings [286] and six-dimensional Myers–Perry black holes [287]. More recently, stable photon orbits have been revealed around Majumdar–Papapetrou di-holes [201, 288], which serve as static axisymmetric toy models for dynamical binary black holes; and around hairy black hole models, including Kerr black holes with scalar hair [184] and Proca hair [79]. In these scenarios, stable photon orbits have been revealed to play a key role in gravitational lensing; in particular, they are associated with distinctive chaotic features in black hole shadows. Non-planar “fundamental photon orbits” (generalisations of light-rings) in generic stationary, axisymmetric spacetimes have been classified – and their stability explored – by Cunha *et al.* [79]. A review of the relationship between fundamental photon orbits and strong-field lensing effects by ultra-compact objects can be found in [60].

In this chapter, we investigate the existence of stable photon orbits in four-dimensional stationary axisymmetric electrovacuum spacetimes; the majority of the work presented in this chapter is based on [289]. In Section 6.2, we review the existence and classification of equatorial circular photon orbits in the charge–spin parameter space for the Kerr–Newman family of spacetimes. In Section 6.3, we employ a Hamiltonian formalism for null geodesics to demonstrate that, for a general stationary axisymmetric (Weyl) spacetime (reviewed in Section 6.3.1), the Einstein–Maxwell equations permit the existence of stable photon orbits in *electrovacuum*, but not in pure vacuum (Section 6.3.2). In Section 6.4, we explore the existence of stable photon orbits in the Reissner–Nordström di-hole class of spacetimes – a subfamily of the general Bretón–Manko–Aguilar two-centre solution [232], which includes the uncharged Weyl–Bach di-hole and extremal Majumdar–Papapetrou di-hole as special cases. Using Poincaré sections, we explore the structure of stable photon orbits around the Majumdar–Papapetrou di-hole in Section 6.5. Finally, we conclude with a discussion of this work and possible extensions in Section 6.6.

6.2 Photon orbits in Kerr–Newman spacetime

6.2.1 Spacetime geometry

Uniqueness and no-hair theorems [40, 290] support the conjecture that the end-product of gravitational collapse in asymptotically flat electrovacuum is a Kerr–Newman black

hole [36, 37]. The most astrophysically relevant case is that of an uncharged Kerr black hole [35], which has become a cornerstone of Einstein’s theory of general relativity. Here, we outline the key features of the Kerr–Newman geometry, paying special attention to its (null) geodesics; for further details, see Sections 2.4 and 2.6, and references therein.

The Kerr–Newman geometry is characterised by three externally observable parameters: mass M , charge Q , and spin $a = \frac{J}{M}$ (where J is the angular momentum). In Boyer–Lindquist coordinates $\{t, r, \theta, \phi\}$, the line element takes the form

$$ds^2 = g_{ab} dx^a dx^b \quad (6.1)$$

$$= - \left(1 - \frac{2Mr}{\Sigma} \right) dt^2 + \frac{\Sigma}{\Delta} dr^2 + \Sigma d\theta^2 + \left(r^2 + a^2 + \frac{2Ma^2r}{\Sigma} \sin^2 \theta \right) \sin^2 \theta d\phi^2 - \frac{4Mar \sin^2 \theta}{\Sigma} dt d\phi, \quad (6.2)$$

where $\Sigma(r, \theta) = r^2 + a^2 \cos^2 \theta$ and $\Delta(r) = r^2 - 2Mr + a^2 + Q^2$. The event horizons are located at Boyer–Lindquist radii $r = r_{\pm} = M \pm \sqrt{M^2 - (a^2 + Q^2)}$. Introducing the charge ratio $q = \frac{Q}{M}$ and spin ratio $\tilde{a} = \frac{a}{M} = \frac{J}{M^2}$, this becomes $r_{\pm} = M \left(1 \pm \sqrt{1 - (\tilde{a}^2 + q^2)} \right)$. The black hole subfamily is given by $\tilde{a}^2 + q^2 \leq 1$, with equality in the extremal case where there is a single event horizon at $r_{\pm} = M$. The Kerr–Newman family may be extended beyond the black hole case to include naked singularities by considering the region of charge–spin parameter space given by the inequality $\tilde{a}^2 + q^2 > 1$.

On Kerr–Newman spacetime, there exist two Killing vector fields, satisfying Killing’s equation $\xi_{(a;b)} = 0$, which are associated with time translations ($\xi_{(t)}^a$), and rotation about the symmetry axis ($\xi_{(\phi)}^a$). These isometries are in one-to-one correspondence with constants of motion along the tangent vector u^a to a geodesic: the energy $E = -\xi_{(t)}^a u_a$, and azimuthal angular momentum $L_z = \xi_{(\phi)}^a u_a$. The existence of a rank-two Killing tensor, which satisfies $K_{(ab;c)} = 0$, is responsible for a further “hidden” constant of motion – the Carter constant, defined as $K = K_{ab} u^a u^b$. The geodesic equations are separable (and hence Liouville integrable) thanks to the existence of this additional conserved quantity; see Section 2.6 for a review.

The radial and latitudinal motion of a massless particle with energy E and angular momentum L_z is described by the equations [176, 291]

$$\Sigma^2 \dot{r}^2 = [E(r^2 + a^2) - aL_z]^2 - K\Delta, \quad (6.3)$$

$$\Sigma^2 \dot{\theta}^2 = K - \left(aE \sin \theta - \frac{L_z}{\sin \theta} \right)^2, \quad (6.4)$$

where an overdot denotes differentiation with respect to an affine parameter along the ray, say λ . Rearrangement of (6.4) permits us to write Carter’s constant in the form

$$K = K_{ab} u^a u^b = \Sigma^2 \dot{\theta}^2 + \left(aE \sin \theta - \frac{L_z}{\sin \theta} \right)^2. \quad (6.5)$$

6.2.2 Classification of equatorial circular photon orbits

We are interested in the problem of classifying equatorial circular photon orbits. Setting $\theta = \frac{\pi}{2}$, we see that (6.5) reduces to $K = (aE - L_z)^2$ and (6.3) gives

$$\dot{r}^2 = \frac{E^2}{r^4} \left[(r^2 + a^2 - ab)^2 - (a - b)^2 (r^2 - 2Mr + a^2 + Q^2) \right], \quad (6.6)$$

where $b = \frac{L_z}{E}$ is the impact parameter. The change of variables $r \mapsto u = \frac{M}{r}$ allows us to express (6.6) as

$$\dot{r}^2 = E^2 \mathcal{R}(u), \quad \mathcal{R}(u) = 1 - (\tilde{b}^2 - \tilde{a}^2) u^2 + (\tilde{a} - \tilde{b})^2 (2u^3 - q^2 u^4), \quad (6.7)$$

where \tilde{a} and q are the spin and charge ratios, respectively, and $\tilde{b} = \frac{b}{M}$. A circular orbit satisfies $\dot{r} = 0 = \ddot{r}$. If $E \neq 0$, the conditions for circular photon orbits can be expressed as $\mathcal{R}(u) = 0 = \mathcal{R}'(u)$, where a prime denotes differentiation with respect to u . Furthermore, an orbit is stable if $\mathcal{R}''(u) < 0$. The problem of classifying equatorial circular photon orbits of non-zero energy therefore reduces to classifying the roots of the quartic $\mathcal{R}(u)$.

We seek the values of the impact parameter b for which \mathcal{R} has repeated roots. Such values may be found by solving $\Delta_u(\mathcal{R}) = 0$, where Δ_u denotes the discriminant with respect to u . The expression $\Delta_u(\mathcal{R})$ is a polynomial in b of degree 10. Phase boundaries in the (q^2, \tilde{a}^2) -plane – which correspond to repeated roots of the polynomial $\Delta_u(\mathcal{R})$ – are obtained by setting “the discriminant of the discriminant” to zero, which factorises as

$$\Delta_b \left(\frac{\Delta_u(\mathcal{R})}{(\tilde{b} - \tilde{a})^6} \right) = 2^{32} (1 - \tilde{a}^2 - q^2) \left[27\tilde{a}^2 - q^2 (9 - 8q^2) \right]^3. \quad (6.8)$$

Hence, the phase boundaries in the charge–spin plane are

$$\tilde{a}^2 = 1 - q^2, \quad (6.9)$$

$$\tilde{a}^2 = \frac{1}{27} q^2 (9 - 8q^2)^2. \quad (6.10)$$

The first of these is the extremality condition for the Kerr–Newman family: solutions with $\tilde{a}^2 + q^2 \leq 1$ are black holes; whereas the converse holds for naked singularities. The phase boundaries intersect at the point $(q, \tilde{a}) = \left(\frac{\sqrt{3}}{2}, \frac{1}{2} \right)$.

In Figure 6.1, we present the Balek–Bičák–Stuchlík [78] phase diagram in the charge–spin plane for equatorial circular photon orbits of positive Boyer–Lindquist radius ($r > 0$) on Kerr–Newman spacetimes. Here we summarise the classification of equatorial circular photon orbits in Regions I–VII of the phase diagram. Regions I and II correspond to the black hole regime with distinct horizons at $r = r_{\pm}$ ($\tilde{a}^2 + q^2 < 1$); Regions III–VIII correspond to naked singularities with no event horizon ($\tilde{a}^2 + q^2 > 1$). All orbits are *unstable*, unless stated otherwise. In Region I there are two exterior orbits ($r > r_+$). In Region II there are two interior orbits ($r < r_-$) and two exterior orbits ($r > r_+$); the innermost orbit is stable. Regions III and IV admit two orbits; the innermost is stable. In Region V there are no equatorial circular photon orbits. Finally, Regions VI and VII admit four orbits; the inner pair is stable. We remark that Regions III and VI (naked singularities with $0 < q^2 < 1$) admit a counter-rotating orbit; whereas Regions IV and VII do not.

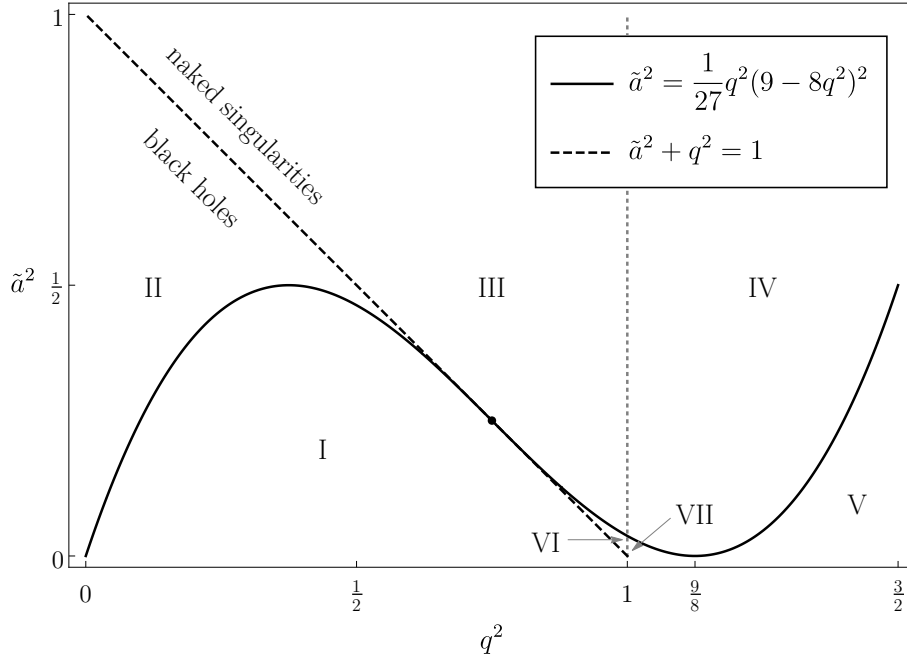


Figure 6.1. Phase diagram for equatorial circular photon orbits of positive Boyer–Lindquist radius ($r > 0$) on Kerr–Newman spacetime in the charge–spin plane (q^2, \tilde{a}^2), where $q = \frac{Q}{M}$ and $\tilde{a} = \frac{a}{M} = \frac{J}{M^2}$. The black solid line shows the phase boundary $\tilde{a}^2 = \frac{1}{27}q^2(9 - 8q^2)^2$; the black dashed line is the phase boundary $\tilde{a}^2 + q^2 = 1$, which separates black holes (Regions I and II) from naked singularities (Regions III–VII). The phase boundaries meet at the point $(q, \tilde{a}) = \left(\frac{\sqrt{3}}{2}, \frac{1}{2}\right)$. Stable equatorial circular photon orbits exist within the inner horizon for black holes in Region II; on (inside) extremal horizons $r = M$ for $a < \frac{1}{2}$ ($a > \frac{1}{2}$); and around naked singularities in Regions III, IV and VI. (Equivalent phase diagrams are presented in Figure 6 of [279], Figure 2.1 of [78], and Figure 4 of [280].)

We now consider the critical cases, which demarcate the seven (open) regions of the (q^2, \tilde{a}^2) -plane considered above. Firstly, for $\tilde{a}^2 = 0$ and $0 \leq q^2 < 1$ (Schwarzschild and Reissner–Nordström black holes), there exists one exterior orbit. For $\tilde{a}^2 = 0$ and $1 \leq q^2 < \frac{9}{8}$ (Reissner–Nordström naked singularities), there are two orbits; the innermost being stable. In the case $q^2 = 0$, $0 < \tilde{a}^2 < 1$ (Kerr black holes), there are two exterior orbits and one interior orbit. (Recall from Section 3.3.3 that the Kerr black hole solution admits a pair of unstable light-rings at $r = \hat{r}_{\pm}$; one of which is co-rotating with the spin of the black hole; the other is counter-rotating. In the Schwarzschild limit $a \rightarrow 0$, these light-rings coincide at $r = 3M$.) In the case of extremal Kerr–Newman black holes, with $\tilde{a}^2 + q^2 = 1$, there are three equatorial circular photon orbits, whose stability depends on the value of \tilde{a} : for $0 < a < \frac{1}{2}$, there is one stable orbit on the horizon ($r = M$) and two exterior orbits; for $\frac{1}{2} < a < 1$, there is one stable interior orbit ($r < M$), one horizon orbit ($r = M$) and one exterior orbit ($r > M$) [282, 283]. At the point $(q, \tilde{a}) = \left(\frac{\sqrt{3}}{2}, \frac{1}{2}\right)$, where the phase boundaries (6.9) and (6.10) intersect, there is a marginally stable orbit ($\mathcal{R}''(u) = 0$) on the extremal horizon, and one exterior orbit at $r = 3M$.

6.3 Photon orbits in stationary axisymmetric spacetimes

6.3.1 Hamiltonian formalism for null geodesics

Stationary axisymmetric spacetime in electrovacuum can be described in Weyl–Lewis–Papapetrou coordinates $\{t, \rho, z, \phi\}$ by the line element [83, 86]

$$ds^2 = g_{ab} dx^a dx^b = -f (dt + w d\phi)^2 + \frac{1}{f} [e^{2\gamma} (d\rho^2 + dz^2) + \rho^2 d\phi^2], \quad (6.11)$$

and electromagnetic one-form potential

$$A_a dx^a = A_t dt + A_\phi d\phi, \quad (6.12)$$

where f , γ , w , A_t and A_ϕ are functions of ρ and z only; see Section 2.3.3 and e.g. [83, 86] for further details. The geodesics $q^a(\lambda)$ are the integral curves of Hamilton's equations, with Hamiltonian function $H(q, p) = \frac{1}{2} g^{ab} p_a p_b$ and canonical momenta $p_a = g_{ab} \dot{q}^a$. Here, an overdot denotes differentiation with respect to the affine parameter λ . The metric components are independent of t and ϕ , so p_t and p_ϕ are constants of motion; one may set $p_t = -1$ without loss of generality, by rescaling the affine parameter. Further, H itself is a constant of motion along geodesics, with $H = 0$ in the null case.

As shown in Section 2.2, null geodesics are invariant under conformal transformations of the form $g_{ab} \mapsto \Omega^2(q) g_{ab}$, where $\Omega > 0$ is a function of the spacetime coordinates. Employing a conformal transformation with $\Omega^2 = e^{-2\gamma} f$, we are able to recast the Hamiltonian associated with (6.11) in two-dimensional ‘‘canonical form’’ as

$$H(\rho, z, p_\rho, p_z; p_\phi) = \frac{1}{2} (p_\rho^2 + p_z^2) + V(\rho, z; p_\phi), \quad (6.13)$$

$$V(\rho, z; p_\phi) = -\frac{e^{2\gamma}}{2\rho^2} \left[\frac{\rho^2}{f^2} - (w + p_\phi)^2 \right], \quad (6.14)$$

where $p_\phi = \text{constant}$ is a free parameter. The positivity of the kinetic term in (6.13) and the Hamiltonian constraint $H = 0$ imply that $V \leq 0$. For a photon with fixed azimuthal angular momentum p_ϕ , this inequality demarcates the *allowed regions* of configuration space.

Null geodesics of constant ρ and z (stationary points) exist where $V = 0$ and $\nabla V = \mathbf{0}$, where the $\nabla = (\partial_\rho, \partial_z)$ is the standard two-dimensional gradient operator. The stability of such orbits can be determined by considering the extrema of the potential V : the orbit is stable if the stationary point is a local minimum of V , i.e., if $\det \mathcal{H}(V) > 0$ and $\text{tr} \mathcal{H}(V) > 0$, where

$$\mathcal{H}(V) = \begin{bmatrix} V_{,\rho\rho} & V_{,\rho z} \\ V_{,\rho z} & V_{,zz} \end{bmatrix} \quad (6.15)$$

is the Hessian matrix for the potential V . In general, null geodesics are kinematically bounded if there exists a contour $V(\rho, z) = 0$ which is closed in the (ρ, z) -plane that bounds an open region in which the potential is negative.

A drawback of the potential (6.14) is that it depends on the orbital parameter p_ϕ . One may factorise the potential as

$$V = -\frac{e^{2\gamma}}{2\rho^2} (h^+ + p_\phi)(h^- - p_\phi), \quad (6.16)$$

where we introduce the *height functions* (or *effective potentials*)

$$h^\pm(\rho, z) = \frac{\rho}{f} \pm w, \quad (6.17)$$

which are independent of the azimuthal angular momentum p_ϕ . (In the case of *static* axisymmetric spacetimes, described by (6.11) with $w = 0$, we write $h^\pm = h$.) In order to determine the closed contours $V = 0$ (which are necessary for kinematically bounded photon orbits), it is sufficient to seek closed contours $h^\pm = \mp p_\phi$. Since $p_\phi \in \mathbb{R}$, closed contours $V = 0$ – and thus stable photon orbits – exist in the neighbourhood of any local *maximum* of the effective potentials (6.17).¹ The problem of classifying stable photon orbits in stationary axisymmetric spacetimes reduces to classifying the fixed points of the effective potentials (6.17), which is achieved by considering $\det \mathcal{H}(h^\pm)$ and $\text{tr} \mathcal{H}(h^\pm)$, where $\mathcal{H}(h^\pm)$ is the Hessian matrix for h^\pm . In particular, a fixed point of h^\pm is a local extremum (saddle point) if $\det \mathcal{H}(h^\pm) > 0$ ($\det \mathcal{H}(h^\pm) < 0$). In the case of local extrema, we have a maximum (minimum) if $\text{tr} \mathcal{H}(h^\pm) < 0$ ($\text{tr} \mathcal{H}(h^\pm) > 0$).

6.3.2 Classification of photon orbits in electrovacuum

We now classify the stationary points of the height function (6.17) for a general stationary axisymmetric geometry in electrovacuum, described by the line element (6.11) and electromagnetic four-potential (6.12). In [85], Ernst presents a formulation of the Einstein–Maxwell field equations for the electrovacuum Weyl solution; the equations relevant for our purposes are

$$(\nabla f)^2 - \rho^{-2} f^4 (\nabla w)^2 + 2f (\nabla A_t)^2 + 2\rho^{-2} f^3 (\nabla A_\phi - w \nabla A_t)^2 = f \nabla^2 f, \quad (6.18)$$

$$\nabla \cdot [\rho^{-2} f^2 \nabla w - 4\rho^{-2} f A_t (\nabla A_\phi - w \nabla A_t)] = 0, \quad (6.19)$$

$$\nabla \cdot [\rho^{-2} f (\nabla A_\phi - w \nabla A_t)] = 0. \quad (6.20)$$

These equations may be obtained by considering the trace-reversed Einstein field equations $R_{ab} = 8\pi \bar{T}_{ab}$, where $\bar{T}_{ab} = T_{ab} - \frac{1}{2} T g_{ab}$, with $T = 0$ for the electromagnetic field (see Section 2.3.2). In (6.18)–(6.18), the two-gradient operator acting on an arbitrary scalar field $\varphi(\rho, z)$ is $\nabla \varphi = (\varphi_{,\rho}, \varphi_{,z})$; the divergence operator acting on an arbitrary vector field $\mathbf{F}(\rho, z) = (F_1(\rho, z), F_2(\rho, z))$ is $\nabla \cdot \mathbf{F} = \frac{1}{\rho}(\rho F_1)_{,\rho} + F_2_{,z}$; and the Laplacian of a scalar field $\varphi(\rho, z)$ is given by $\nabla^2 \varphi = \nabla \cdot (\nabla \varphi) = \varphi_{,\rho\rho} + \frac{1}{\rho} \varphi_{,\rho} + \varphi_{,zz}$.

To classify the stationary points of h^\pm , we seek an expression for $\text{tr} \mathcal{H}(h^\pm) = h^\pm_{,\rho\rho} + h^\pm_{,zz}$. Taking second derivatives of (6.17), we find

$$\text{tr} \mathcal{H}(h^\pm) = -\rho f^{-3} (f \nabla^2 f) - f^{-2} f_{,\rho} + 2\rho f^{-3} (\nabla f)^2 \pm (w_{,\rho\rho} + w_{,zz}). \quad (6.21)$$

The field equation (6.19) rearranges to

$$w_{,\rho\rho} + w_{,zz} = \rho^{-1} f^2 w_{,\rho} - 2f^{-1} \nabla f \cdot \nabla w + 4\rho^2 f^{-2} \nabla \cdot [\rho^{-2} f A_t (\nabla A_\phi - w \nabla A_t)]. \quad (6.22)$$

Inserting (6.18) and (6.22) into the right-hand side of (6.21) yields

$$\begin{aligned} \text{tr} \mathcal{H}(h^\pm) &= \rho f^{-3} (\nabla f)^2 \mp 2f^{-1} \nabla f \cdot \nabla w + \rho^{-1} f (\nabla w)^2 - f^{-2} f_{,\rho} \pm \rho^{-1} w_{,\rho} \\ &\quad - 2\rho f^{-2} (\nabla A_t)^2 - 2\rho^{-1} (\nabla A_\phi - w \nabla A_t)^2 \\ &\quad \pm 4\rho^2 f^{-2} \nabla \cdot [\rho^{-2} f A_t (\nabla A_\phi - w \nabla A_t)]. \end{aligned} \quad (6.23)$$

¹Note that local maxima of h^\pm correspond to local minima of V , due to the minus sign in (6.16).

Suppose we have found a stationary point of either h^+ or h^- , that is, a point (ρ, z) which satisfies the stationary point conditions $h^{\pm}_{,\rho} = 0 = h^{\pm}_{,z}$. Taking partial derivatives of (6.17) and employing the stationary point conditions, we see that the equations

$$w_{,\rho} = \pm (\rho f^{-2} f_{,\rho} - f^{-1}), \quad (6.24)$$

$$w_{,z} = \pm \rho f^{-2} f_{,z} \quad (6.25)$$

must hold at fixed points of h^{\pm} . One can use (6.24) and (6.25) to eliminate terms involving partial derivatives of w from the right-hand side of (6.23); at stationary points of h^{\pm} , the relevant terms can be written

$$2f^{-1} \nabla f \cdot \nabla w = \pm \left[2\rho f^{-3} (\nabla f)^2 - 2f^{-2} f_{,\rho} \right], \quad (6.26)$$

$$\rho^{-1} f (\nabla w)^2 = \rho f^{-3} (\nabla f)^2 - 2f^{-2} f_{,\rho} + \rho^{-1} f^{-1}. \quad (6.27)$$

Now, substituting (6.24), (6.26) and (6.27) into the right-hand side of (6.23), we see that the first five terms vanish; this leaves us with

$$\begin{aligned} \text{tr } \mathcal{H}(h^{\pm}) &= -2\rho f^{-2} (\nabla A_t)^2 - 2\rho^{-1} (\nabla A_{\phi} - w \nabla A_t)^2 \\ &\pm 4\rho^2 f^{-2} \nabla \cdot [\rho^{-2} f A_t (\nabla A_{\phi} - w \nabla A_t)]. \end{aligned} \quad (6.28)$$

Consider the final term on the right-hand side of (6.28). Expanding this term using the product rule for the divergence operator, we find

$$\begin{aligned} 4\rho^2 f^{-2} \nabla \cdot [\rho^{-2} f A_t (\nabla A_{\phi} - w \nabla A_t)] &= 4f^{-1} \nabla A_t \cdot (\nabla A_{\phi} - w \nabla A_t) \\ &+ 4\rho^2 f^{-2} A_t \nabla \cdot [\rho^{-2} f (\nabla A_{\phi} - w \nabla A_t)]. \end{aligned} \quad (6.29)$$

Availing the remaining field equation (6.20), we see that the final term on the right-hand side of (6.29) vanishes, such that

$$4\rho^2 f^{-2} \nabla \cdot [\rho^{-2} f A_t (\nabla A_{\phi} - w \nabla A_t)] = 4f^{-1} \nabla A_t \cdot (\nabla A_{\phi} - w \nabla A_t). \quad (6.30)$$

Finally, inserting (6.30) into the right-hand side of (6.28) yields

$$\begin{aligned} \text{tr } \mathcal{H}(h^{\pm}) &= -2\rho f^{-2} (\nabla A_t)^2 - 2\rho^{-1} (\nabla A_{\phi} - w \nabla A_t)^2 \\ &\pm 4f^{-1} \nabla A_t \cdot (\nabla A_{\phi} - w \nabla A_t), \end{aligned} \quad (6.31)$$

which factorises to give

$$\text{tr } \mathcal{H}(h^{\pm}) = -\frac{2}{\rho} |\mathbf{W}_{\pm}|^2, \quad \mathbf{W}_{\pm} = h^{\pm} \nabla A_t \mp \nabla A_{\phi}, \quad (6.32)$$

where $|\mathbf{W}_{\pm}|^2 = \mathbf{W}_{\pm} \cdot \mathbf{W}_{\pm}$. Using (6.32), we can replace the terms involving $h^{\pm}_{,zz}$ in the expression for $\det \mathcal{H}(h^{\pm})$; this gives

$$\det \mathcal{H}(h^{\pm}) = - \left[(h^{\pm}_{,\rho\rho})^2 + \frac{2}{\rho} |\mathbf{W}_{\pm}|^2 h^{\pm}_{,\rho\rho} + (h^{\pm}_{,\rho z})^2 \right]. \quad (6.33)$$

In the pure vacuum case, the electromagnetic four-potential (6.12) is zero, so $\mathbf{W}_{\pm} = 0$ and the right-hand side of (6.32) vanishes, i.e., $\text{tr } \mathcal{H}(h^{\pm}) = 0$. Moreover, $\det \mathcal{H}(h^{\pm}) =$

$-\left[(h^\pm_{,\rho\rho})^2 + (h^\pm_{,\rho z})^2\right] \leq 0$. Hence, h^\pm cannot possess a first-order local maximum, so generic stable photon orbits are ruled out in pure vacuum.

For $\nabla A_a \neq \mathbf{0}$, we have $\mathbf{W}_\pm \neq \mathbf{0}$, so $\text{tr } \mathcal{H}(h^\pm) < 0$ by (6.32). Therefore local minima of h^\pm are forbidden; however, local maxima cannot be ruled out. Local maxima of h^\pm , and hence stable photon orbits, are possible if the right-hand side of (6.33) is positive; this can only be true if $h^\pm_{,\rho\rho} < 0$ and $(h^\pm_{,\rho z})^2 < -h^\pm_{,\rho\rho} \left(h^\pm_{,\rho\rho} + \frac{2}{\rho} |\mathbf{W}_\pm|^2\right)$. In the special case of height functions with an equatorial symmetry, the derivative $h^\pm_{,\rho z}$ vanishes, and the stability conditions reduce to the inequality $0 < -h^\pm_{,\rho\rho} < \frac{2}{\rho} |\mathbf{W}_\pm|^2$.

The absence of stable photon orbits in static axisymmetric *vacuum* spacetimes has been pointed out by Liang [77], who remarked that “stability within the plane implies instability off the plane”. Equation (6.32) demonstrates how Liang’s result breaks down in *electrovacuum*, when an electromagnetic field is introduced. Generalisations of the results presented in this section – to include other fields, matter sources, or a non-zero cosmological constant – may be considered, and would provide a more complete description of the geodesic dynamics on stationary axisymmetric spacetimes in Einstein’s general theory of relativity. This is discussed further in Section 6.6.1.

6.4 Existence of stable photon orbits around di-holes

6.4.1 Di-hole formalism: the Bretón–Manko–Aguilar solution

We consider the $N = 2$ Bretón–Manko–Aguilar family of electrostatic solutions to the Einstein–Maxwell equations [232, 292–294]. In Appendix D, we briefly review the so-called “physical parametrisation” of the Bretón–Manko–Aguilar solution, first presented by Manko [295]. In this parametrisation, the Bretón–Manko–Aguilar family of solutions is described by five parameters: the black hole masses M_\pm ; the black hole charges Q_\pm ; and the separation of centres d .

Here, we restrict our attention to the Reissner–Nordström di-hole subfamily, in which each of the sources is a black hole (i.e., $Q_\pm \leq M_\pm$, with equality in the case of extremal black holes); we do not consider naked singularities ($Q_\pm > M_\pm$). The Reissner–Nordström di-hole family contains as special cases (i) the Majumdar–Papapetrou di-hole family [202, 203], in which both black holes are extremal ($Q_\pm = M_\pm$); (ii) the Weyl–Bach di-hole family [222], in which both black holes are uncharged ($Q_\pm = 0$). With the exception of the Majumdar–Papapetrou subfamily, the black holes which comprise the Reissner–Nordström di-hole are held in equilibrium by a “Weyl strut” [222], and their (outer) event horizons appear as “rods” on the symmetry axis of coordinate length $2\sqrt{M_\pm^2 - Q_\pm^2}$ in Weyl–Lewis–Papapetrou coordinates. In the Majumdar–Papapetrou case ($Q_\pm = M_\pm$), the black hole horizons appear as points (of finite proper area) on the z -axis.

In this section we consider two special cases of the Reissner–Nordström di-hole family: (i) the Majumdar–Papapetrou di-hole family, characterised by the mass ratio parameter $\eta = \frac{M_+ - M_-}{M_+ + M_-}$ (Section 6.4.2); (ii) the equal-mass, equal-charge RN di-hole, characterised by the charge ratio parameter $q = \frac{Q_+}{M_+}$ (Section 6.4.3). We analyse the existence of stable photon orbits in both cases; this involves employing a numerical root finder to classify the

fixed points of the static height function h , given by (6.17) with $w = 0$.

6.4.2 Majumdar–Papapetrou di-hole

The Majumdar–Papapetrou di-hole comprises a pair of extremal Reissner–Nordström black holes in static equilibrium [89, 202, 203]. The geometry is described by the Weyl line element (6.11) with $\gamma = 0 = w$, and $f = \frac{1}{U^2}$, where

$$U(\rho, z) = 1 + \frac{M_+}{\sqrt{\rho^2 + (z - z_+)^2}} + \frac{M_-}{\sqrt{\rho^2 + (z - z_-)^2}}. \quad (6.34)$$

The only non-zero component of the electromagnetic four-potential (6.12) is $A_t = \frac{1}{U}$. We consider the general case, in which the black hole masses are given by $M_{\pm} = \frac{1}{2}(1 \pm \eta)M_0$, where $\eta \in (-1, 1)$ parametrises the mass ratio, and $M_0 = M_+ + M_-$ is the total mass of the system. The black holes are located along the z -axis ($\rho = 0$) at $z_{\pm} = \pm \frac{dM_{\pm}}{M_0}$, where d denotes the separation between the sources (see Section 4.2.1). The solution is a special case of the Bretón–Manko–Aguilar family (see Section 6.4.1), where each black hole is extremal ($Q_{\pm} = M_{\pm}$). The Majumdar–Papapetrou di-hole was used as a surrogate model to investigate the qualitative features of binary black hole shadows in Chapters 4 and 5.

For the Majumdar–Papapetrou di-hole, the height function takes the form $h = \rho U^2$. As described in Section 6.3, the stationary points of h correspond to orbits of constant ρ and z ; such orbits are stable if the stationary point is a local maximum of h . Moreover, the contours of h around a local maximum will be closed curves, permitting kinematically bounded orbits. We now highlight the existence of stable photon orbits around the Majumdar–Papapetrou di-hole with mass parameter η .

First, consider the equal-mass case, in which $\eta = 0$. We write $M = M_{\pm}$ for the mass of the black holes. The system admits stable photon orbits for separations in the range $\sqrt{\frac{16}{27}} < \frac{d}{M} < \sqrt{\frac{32}{27}}$; this existence region can be determined using the result on stable photon orbits derived in Section 6.3.2, or an alternative approach; both methods are presented in Appendix B. In this parameter regime, h admits four stationary points: a local maximum in the equatorial plane at $(\rho, z) = (\rho_{(1)}, 0)$, with $h = p_{\phi}^{(1)}$; a saddle point in the equatorial plane at $(\rho, z) = (\rho_{(2)}, 0)$, with $\rho_{(2)} > \rho_{(1)}$ and $h = p_{\phi}^{(2)} < p_{\phi}^{(1)}$; and two saddle points out of the plane at $(\rho, z) = (\rho_{(3)}, \pm z_{(3)})$, with $\rho_{(3)} < \rho_{(1)}$ and $h = p_{\phi}^{(3)} < p_{\phi}^{(1)}$. Figures 6.2(a)–6.2(c) show the stationary points and critical contours of the height function h in the (ρ, z) -plane, for a selection of values of d (close to the highly symmetric case $d = M$).

For separations in the range $\sqrt{\frac{16}{27}}M < d < M$, we have $p_{\phi}^{(2)} < p_{\phi}^{(3)}$. Null geodesics with $p_{\phi}^{(2)} \leq p_{\phi} < p_{\phi}^{(3)}$ are permitted to fall into the black holes, but are not able to escape to infinity. Moreover, rays with $p_{\phi}^{(3)} \leq p_{\phi} < p_{\phi}^{(1)}$ which start in the vicinity of the local maximum are kinematically bounded. This is illustrated Figure 6.2(a), where we set $d = 0.95M$.

On the other hand, for $M < d < \sqrt{\frac{32}{27}}M$, we have $p_{\phi}^{(3)} < p_{\phi}^{(2)}$, so rays with $p_{\phi}^{(3)} \leq p_{\phi} < p_{\phi}^{(2)}$ are connected to infinity but not to the black holes. Null rays which start close to the local maximum of h with $p_{\phi}^{(2)} \leq p_{\phi} < p_{\phi}^{(1)}$ are kinematically bounded. An example of this case is shown in Figure 6.2(c), in which $d = 1.05M$.

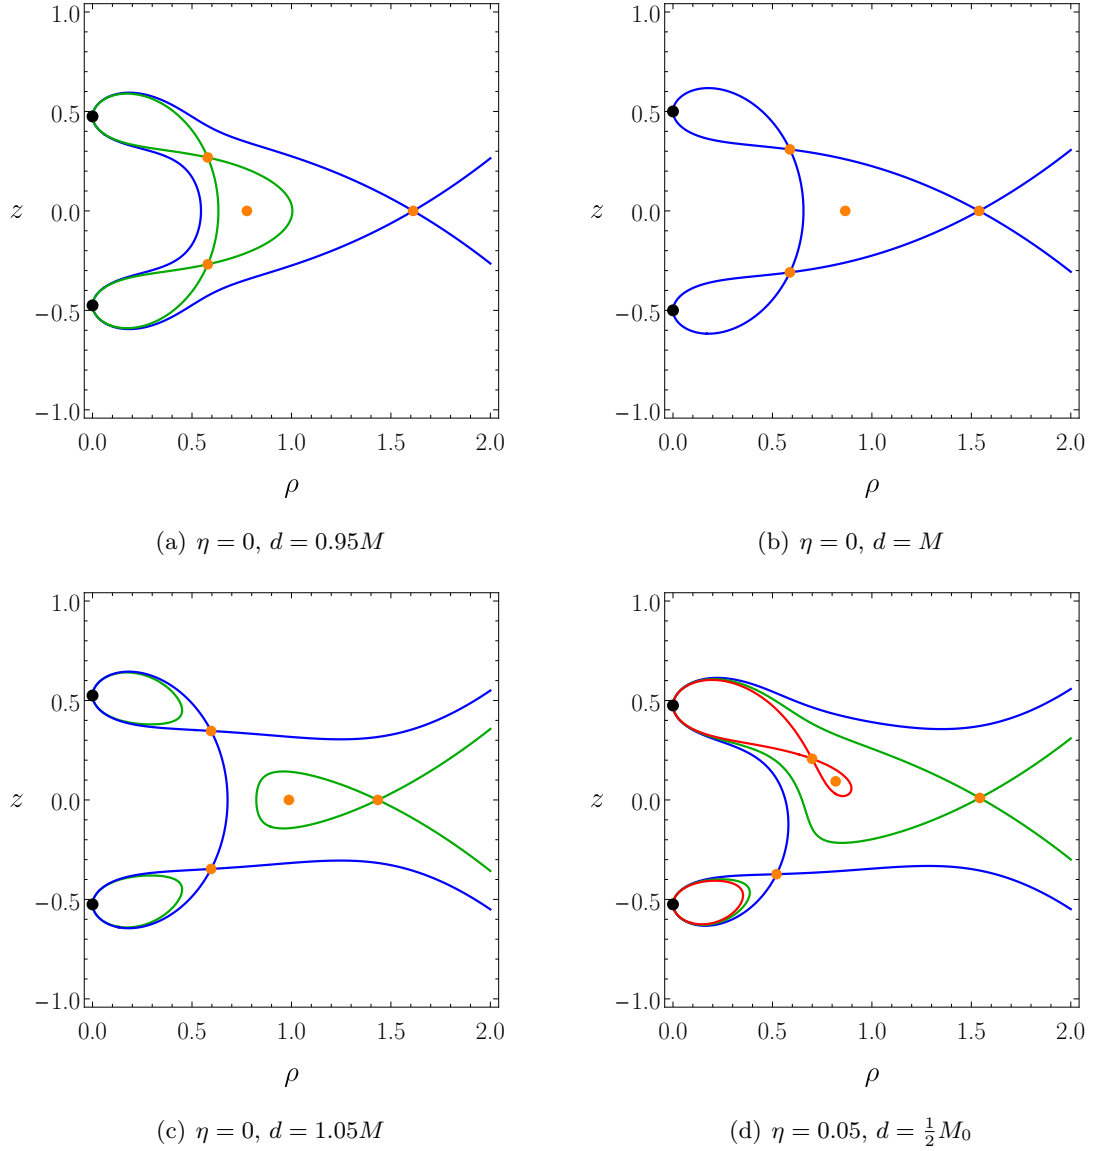


Figure 6.2. Critical contours of the effective potential h for the Majumdar–Papapetrou di-hole, characterised by mass ratio η and separation parameter d . Filled orange circles indicate stationary points of h : three saddle points and one local maximum. Filled black circles on the z -axis represent the black hole horizons, which are located at $z = z_{\pm}$. (a)–(c) Equal-mass di-holes with varying separation d . (d) An unequal-mass di-hole with mass parameter $\eta = 0.05$, and separation of centres $d = \frac{1}{2}M_0$.

In the critical case $d = M$, we have $p_\phi^{(2)} = p_\phi^{(3)}$, and we are able to find closed-form expressions for the fixed points of h and the critical contours which pass through them (see Appendix B). There exists a local maximum at $\rho_{(1)} = \frac{\sqrt{3}}{2}M$, with $p_\phi^{(1)} = \frac{9\sqrt{3}}{2}M$. The maximum is enclosed by a single closed critical contour $h = p_\phi^{(2)} = p_\phi^{(3)} = \frac{1}{2}5^{5/4}\varphi^{3/2}M$, which connects three saddles at $\rho_{(2)} = \frac{1}{2}5^{1/4}\varphi^{3/2}M$, $z = 0$ and $\rho_{(3)} = \frac{1}{2}5^{1/4}\varphi^{-1/2}M$, $\pm z_{(3)} = \pm \frac{M}{2\varphi}$, where $\varphi = \frac{1}{2}(1 + \sqrt{5})$ denotes the golden ratio. This special case is illustrated in Figure 6.2(b).

Figures 4.13(a) and 4.13(c) of Section 4.7 show the critical contours of the effective potential $h = \rho U^2$ in the extreme cases $d = \sqrt{\frac{16}{27}}M$ and $d = \sqrt{\frac{32}{27}}M$, respectively. When $d = \sqrt{\frac{16}{27}}M$, h admits two saddle points in the equatorial plane. For $d = \sqrt{\frac{32}{27}}M$, there are two non-equatorial saddle points and there is a ‘‘cusp’’ in the plane; decreasing d slightly from this marginal value, the cusp becomes a saddle point in the equatorial plane, as shown in Figure 4.13(b). For separations in the range $\sqrt{\frac{16}{27}}M < d < \sqrt{\frac{32}{27}}M$, the effective potential admits one equatorial saddle point, two non-equatorial saddles, and a local maximum in the equatorial plane. The latter corresponds to a stable circular photon orbit.

In the more general case of unequal-mass black holes ($\eta \neq 0$), the \mathbb{Z}_2 symmetry in the equatorial plane is broken: the non-planar saddle points are no longer connected by the same critical contour, and the kinematically bounded ‘‘stable photon orbit region’’ is pulled towards the more massive black hole. This is illustrated in Figure 6.2(d), where the mass ratio is taken to be $\eta = \frac{M_+ - M_-}{M_+ + M_-} = 0.05$ ($M_+ > M_-$), and the separation between the centres is $d = \frac{1}{2}M_0 = \frac{1}{2}(M_+ + M_-)$.

We now consider the effect of increasing the mass ratio on the existence of stable photon orbits around Majumdar–Papapetrou di-holes. In general, we are not able to find closed-form expressions for the stable photon orbit existence regions. For a fixed value of η , we employ a numerical root finder to solve the stationary point conditions $h_{,\rho} = 0 = h_{,z}$, and the stability condition $\det \mathcal{H}(h) = h_{,\rho\rho} h_{,zz} - (h_{,\rho z})^2 = 0$, in order to find the critical values of the dimensionless separation parameter $\tilde{d} = \frac{2d}{M_0}$ for which stable photon orbits exist. In Figure 6.3(a), we present a plot of the (η, \tilde{d}) parameter space which highlights the existence region for stable photon orbits. We observe that increasing the mass ratio has the effect of diminishing the stable photon orbit existence region; for $\eta \gtrsim 0.13$ stable photon orbits do not exist. When $\eta = 0$ (equal-mass case), we see that the existence region is given by $\sqrt{\frac{16}{27}}M < d < \sqrt{\frac{32}{27}}M$, which agrees with the analytical results described above.

6.4.3 Equal-mass, equal-charge Reissner–Nordström di-hole

We now turn our attention to the equal-mass, equal-charge Reissner–Nordström di-hole. This is a two-parameter subfamily of the five-parameter Bretón–Manko–Aguilar solution, presented in Section 6.4.1. This solution comprises a pair of equal-mass black holes ($M_\pm = M$), with charges $Q_\pm = qM$, where $q \in [0, 1]$ is the charge ratio. The case $q = 0$ corresponds to the uncharged Weyl–Bach di-hole, whilst the extremal $q = 1$ case corresponds to the equal-mass Majumdar–Papapetrou di-hole (described in Section 6.4.2).

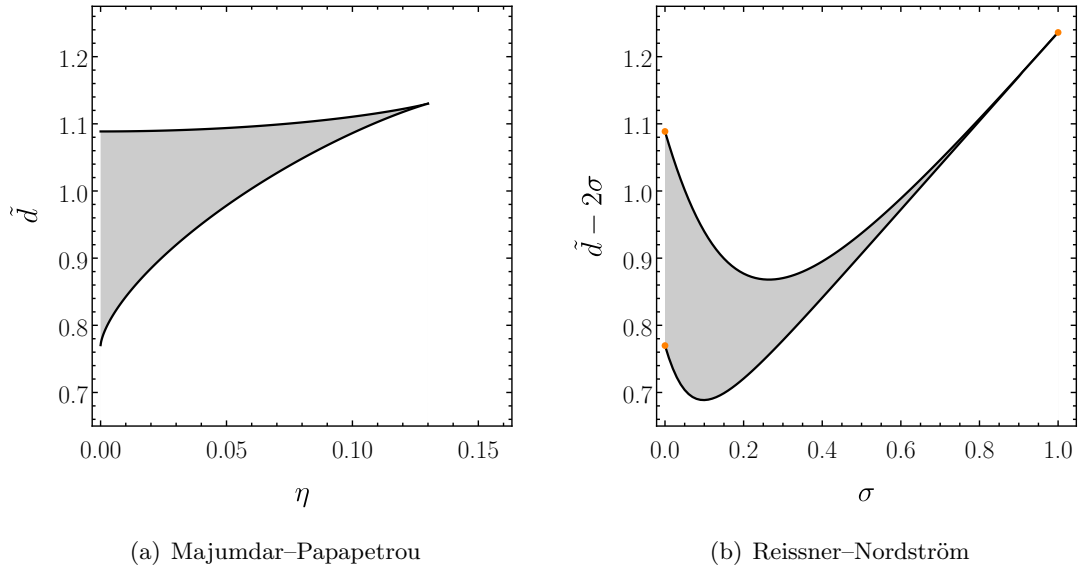


Figure 6.3. Existence region for stable photon orbits around Reissner–Nordström di-holes, a subfamily of the $N = 2$ Bretón–Manko–Aguilar class of electrostatic solutions. (a) Majumdar–Papapetrou di-holes with mass ratio η and separation parameter $\tilde{d} = \frac{2d}{M_0}$. (b) Equal-mass Reissner–Nordström di-holes with charge parameter $\sigma = \sqrt{1 - q^2}$; the separation between the event horizons is $d - 2M\sigma$.

The centres are held in static equilibrium by a Weyl strut, which imparts a force [295]

$$\mathcal{F} = \frac{M\sigma^2}{d^2 - 4M^2\sigma^2}, \quad (6.35)$$

where $\sigma = \sqrt{1 - q^2}$ is a charge parameter. In Weyl–Lewis–Papapetrou coordinates, the black hole event horizons appear as “rods” of length $2\sqrt{M_{\pm}^2 - Q_{\pm}^2} = 2\sigma$ on the z -axis; the coordinate distance between the horizons is given by $d - 2M\sigma$.

Employing a similar method to that described in Section 6.4.2, we consider the effect of varying the charge parameter σ on the existence region for stable photon orbits in the equal-mass, equal-charge Reissner–Nordström di-hole spacetime. The results of the numerical algorithm are presented in Figure 6.3(b). We observe that stable photon orbits exist for $0 \leq \sigma < 1$, i.e., up to but not including the uncharged Weyl–Bach limit ($\sigma = 1$) [222]. (Recall from our key result of Section 6.3.2 that stable photon orbits are permitted in electrovacuum, but not in pure vacuum.) Increasing the charge parameter σ has the effect of shrinking the existence region for stable photon orbits. In Figure 6.3(b), the critical values of the separation parameter for the Majumdar–Papapetrou and Weyl–Bach cases are shown as orange points: the analytical values are $d = \sqrt{\frac{16}{27}}M$ and $d = \sqrt{\frac{32}{27}}M$ (Majumdar–Papapetrou); and $d = 2\left(\frac{1}{\varphi} + 1\right)M$ (Weyl–Bach), where φ denotes the golden ratio. (The Weyl–Bach case was studied by Coelho and Herdeiro in [223].)

6.5 Geodesic structure of stable photon orbits

In general, the two-dimensional Hamiltonian system given by the Hamiltonian function (6.13) with potential (6.14) is non-integrable. (The stationary axisymmetric geometry (6.11) will not admit higher-order Killing tensor, except in special cases.) Thus, stable photon orbits which exist close to local maxima of the effective potential(s) (6.17) will exhibit rich – possibly chaotic – behaviour. In this section, we analyse the geodesic dynamics of stable photon orbits for the Majumdar–Papapetrou di-hole, described in Section 6.4.2.

Introducing a dimensionless parameter μ , we express the azimuthal angular momentum in the form

$$p_\phi = (1 - \mu) p_\phi^* + \mu p_\phi^{(1)}, \quad p_\phi^* = \max(p_\phi^{(2)}, p_\phi^{(3)}). \quad (6.36)$$

The contours of h form closed curves on a subset of the (ρ, z) -plane for $0 \leq \mu < 1$. Orbits which start in this region are confined to a toroidal region in the three-dimensional space spanned by the coordinates $\{\rho, z, \phi\}$; such orbits are forbidden from escaping by angular momentum. Thus, kinematically bounded stable photon orbits exist in the regime $0 \leq \mu < 1$.

Close to the local maximum of h (i.e., in the limit $\mu \rightarrow 1$ from below), the bounded region forms a small ellipse in the (ρ, z) -plane. At its local minimum $(\rho, z) = (\rho_0, z_0)$, the potential (6.14) satisfies $V(\rho_0, z_0) = 0$, and $V_{,\rho}(\rho_0, z_0) = 0 = V_{,z}(\rho_0, z_0)$. Performing a second-order Taylor expansion about this minimum, the potential may be expressed as

$$V \sim \frac{1}{2} \left[V_{,\rho\rho} (\rho - \rho_0)^2 + 2V_{,\rho z} (\rho - \rho_0)(z - z_0) + V_{,zz} (z - z_0)^2 \right]. \quad (6.37)$$

One can then perform a rotation in the (ρ, z) -plane to remove the mixed quadratic term in the potential (6.37). The resulting potential resembles that of an anisotropic harmonic oscillator $V = \frac{1}{2} \left[\omega_\rho^2 (\rho - \rho_0)^2 + \omega_z^2 (z - z_0)^2 \right]$, with characteristic frequencies given by $\omega_\rho = \sqrt{V_{,\rho\rho}}$ and $\omega_z = \sqrt{V_{,zz}}$. For $d < M$ ($d > M$), we have $\omega_\rho > \omega_z$ ($\omega_\rho < \omega_z$). The case $d = M$ is *isotropic*, in the sense that $\omega_\rho = \omega_z$.

From numerical investigations, we see that generic kinematically bounded photon orbits for $\mu \rightarrow 0$ are precessing ellipses, confined to a small elliptical region of the (ρ, z) -plane in the neighbourhood of the local maximum of the effective potential. Even in this regime, higher-order corrections to the potential V cannot be neglected if the characteristic frequencies of the orbit are commensurate, i.e., $\frac{\omega_\rho}{\omega_z} \approx \frac{n_1}{n_2}$, where n_1 and n_2 are small positive integers. This is because perturbation methods generally break down due to secular terms which appear in the case of low-order resonances [142]. Interestingly, the isotropic case admits a 1:1 resonance, and other low-order resonances.

6.5.1 Bounded orbits and Poincaré sections

The dynamics of kinematically bounded photon orbits can be analysed using Poincaré sections. For a review of this method, see Section 2.6.5 and [103], for example. In Appendix C, we present a brief, illustrative review of Poincaré sections for the well-known Hénon–Heiles Hamiltonian system [117]. A Poincaré map allows us to build up a picture of the

global dynamics, which is governed by a Hamiltonian flow on a three-dimensional energy hypersurface (here $H = 0$, with $p_\phi = \text{constant}$), by constructing a discrete dynamical system on a two-dimensional subspace – the Poincaré section.

We let $\mathcal{E} = \{(\rho, z, p_\rho, p_z) \mid H = 0\}$ be the three-dimensional energy hypersurface, defined by the Hamiltonian constraint. We take the two-dimensional Poincaré section (or surface of section) to be the surface $\mathcal{S} = \{(\rho, p_\rho) \mid z = 0\}$, which intersects \mathcal{E} . The Poincaré map $P: \mathcal{S} \rightarrow \mathcal{S}$ then maps a trajectory on \mathcal{E} that begins on \mathcal{S} to its next point of intersection with \mathcal{S} .

Here, we construct Poincaré sections for the equal-mass Majumdar–Papapetrou di-hole using a method analogous to that of the Hénon–Heiles system presented in Appendix C. We consider a selection of orbits which start within the closed contour $h = p_\phi$, where $p_\phi^* \leq p_\phi < p_\phi^{(1)}$. The trajectories are numerically integrated using Hamilton’s equations. Each time the orbit intersects the surface $z = 0$ (with $p_z > 0$), the values of ρ and p_ρ are recorded.

Figures 6.4 and 6.5 show Poincaré sections for the isotropic case $d = M$, where we consider the effect of varying the dimensionless angular momentum parameter μ , defined in (6.36). We also present selected examples of kinematically bounded orbits projected onto the (ρ, z) -plane. The orbits fall into three categories: (i) low-order $n_1 : n_2$ resonant orbits; (ii) high-order resonances (or box orbits); and (iii) chaotic orbits. (See [296] for a classification of these types of orbits for the Hénon–Heiles system.) The $n_1 : n_2$ families of resonant orbits (where n_1 and n_2 are both small positive integers) exhibit an oscillatory pattern: the orbit completes n_1 oscillations in the ρ -direction in the time it takes to complete n_2 oscillations in the z -direction. The ratio $n_1 : n_2$ corresponds to the ratio of the characteristic frequencies of the orbit, which are defined with respect to the frequency of the largest amplitude in each direction; see [296] and references therein.

For $\mu = 0.75$, the Poincaré section, shown in Figure 6.4(a), exhibits four regular regions, delineated by a “separatrix”. Each of these four regions contains closed, smooth, oval-shaped curves, which surround stable elliptic fixed points of the Poincaré map. The pair of fixed points in the plane $p_\rho = 0$ are “rotational” orbits which form closed curves in the (ρ, z) -plane; whereas the fixed points out of the plane are “librational” orbits, which correspond to open curves in the (ρ, z) -plane that bounce between two sides of the closed contour $h = p_\phi$. The separatrix intersects itself three times; the points of intersection are unstable hyperbolic fixed points of the Poincaré map. The orbits corresponding to the elliptic and hyperbolic fixed points are shown in Figure 6.4(b): the stable rotational orbit is shown in red; the stable librational orbits are shown in green; and the unstable orbits are shown in magenta. In Figure 6.4(c), we show an orbit whose initial conditions are perturbed slightly from those of the planar unstable hyperbolic fixed point. Figure 6.4(d) shows an example of a box orbit (i.e., a higher-order resonance).

We now turn our attention to Figure 6.5, in which we examine the effect of decreasing the parameter μ (i.e., approaching the threshold value $p_\phi = p_\phi^*$). In Figure 6.5(a), we show the Poincaré section for $\mu = 0.4$. The four regions around the elliptic fixed points containing regular orbits persist; however, the separatrix has degenerated into a chaotic band. Moreover, stable and unstable Kolmogorov–Arnold–Moser (KAM) islands

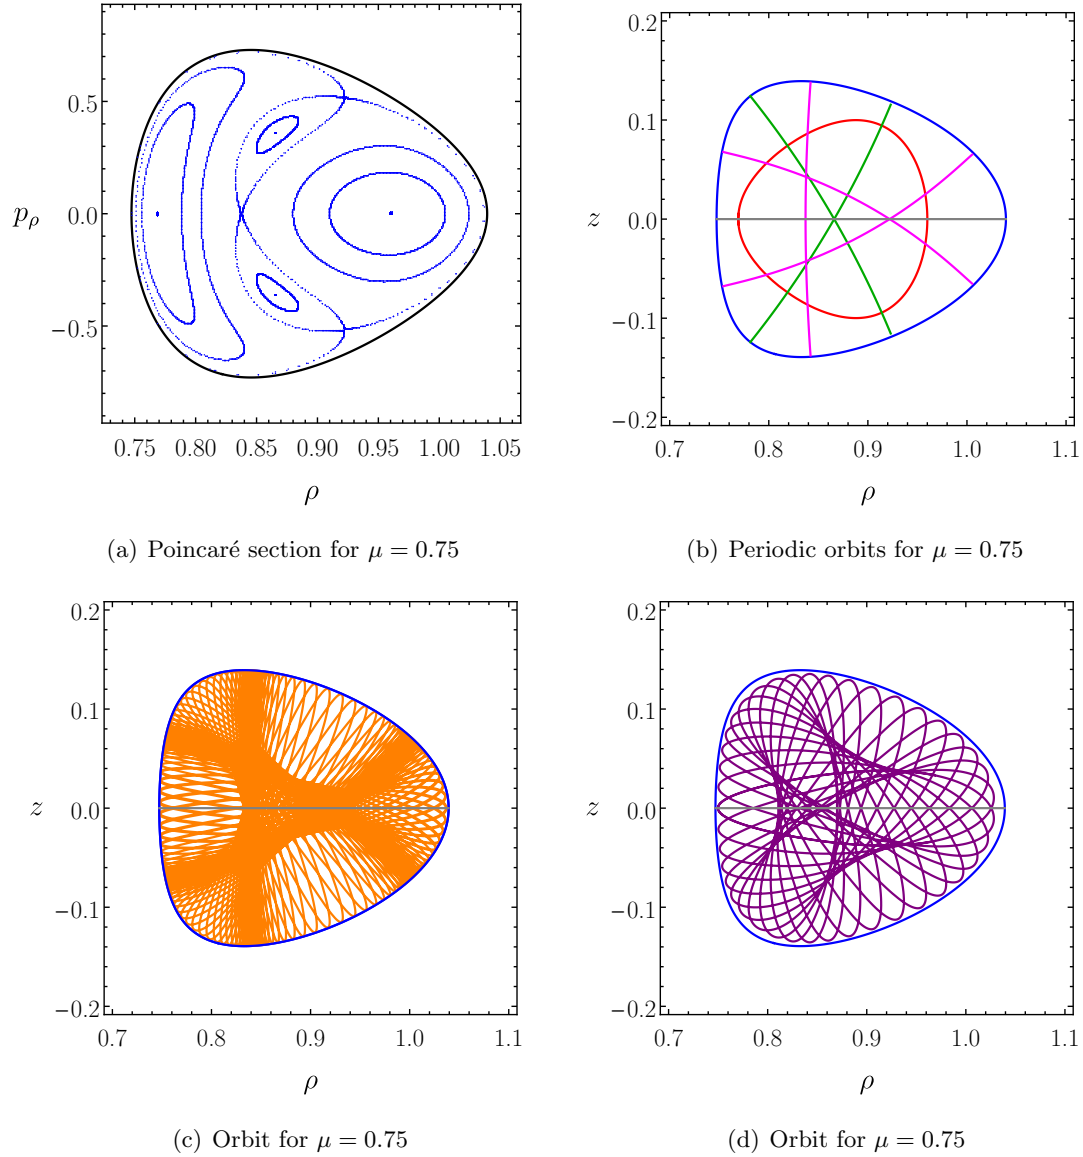


Figure 6.4. (a) Poincaré section in the (ρ, p_ρ) -plane for a selection of kinematically bounded photon orbits with dimensionless angular momentum parameter $\mu = 0.75$. (b) Periodic orbits projected onto the (ρ, z) -plane: a rotational 1 : 1 resonant orbit (“loop orbit”) [red], and a pair of librational 1 : 1 resonant orbits (“linear orbits”) [green], corresponding to elliptic fixed points of the Poincaré map; three unstable linear orbits [magenta], corresponding to hyperbolic fixed points. (c) An orbit whose initial conditions are perturbed slightly from those of the planar hyperbolic fixed point. (d) A high-order resonance (“box orbit”).

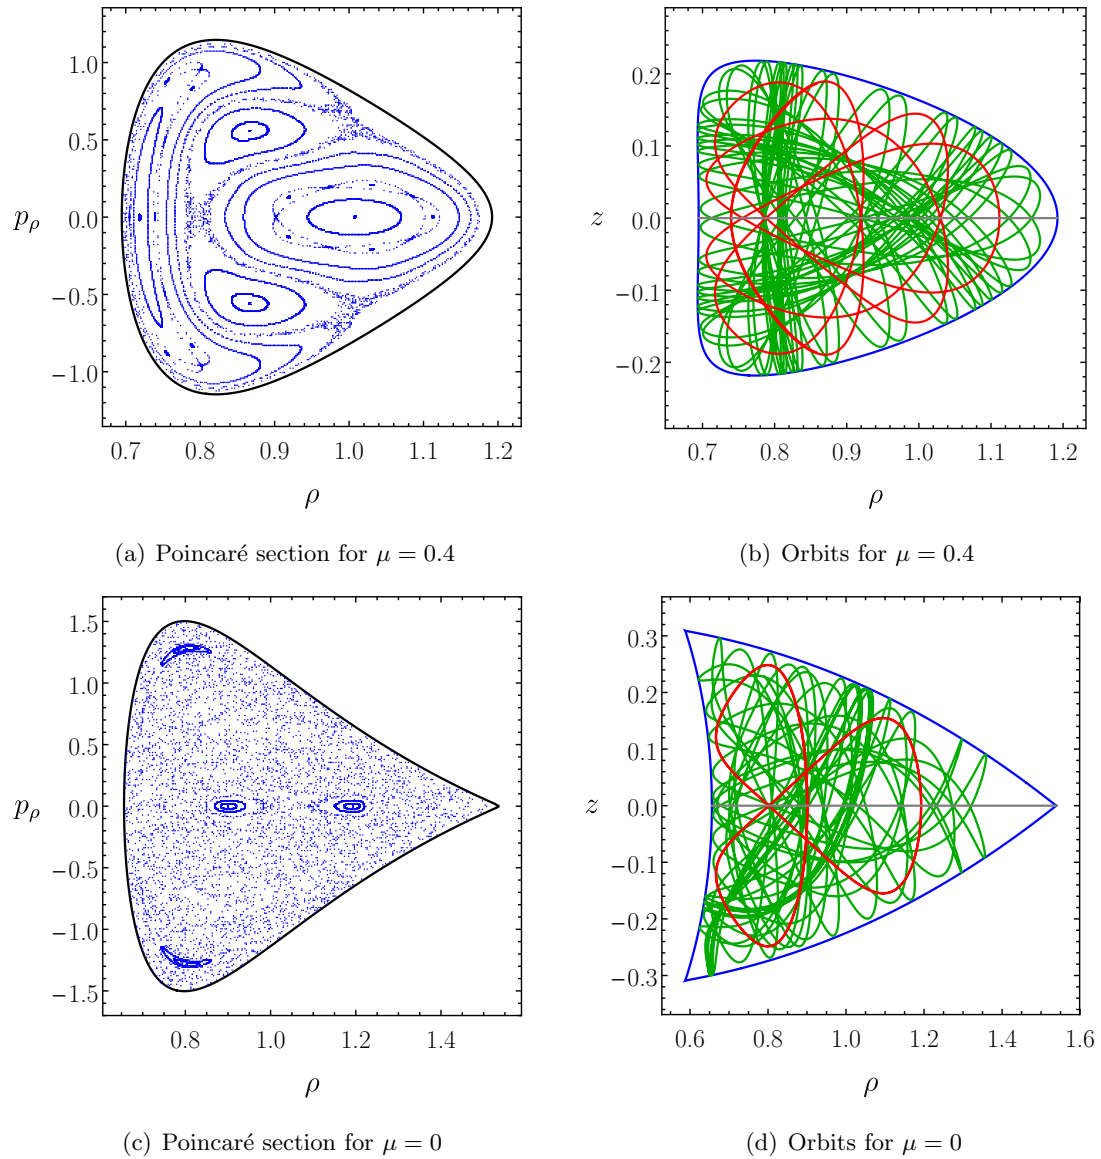

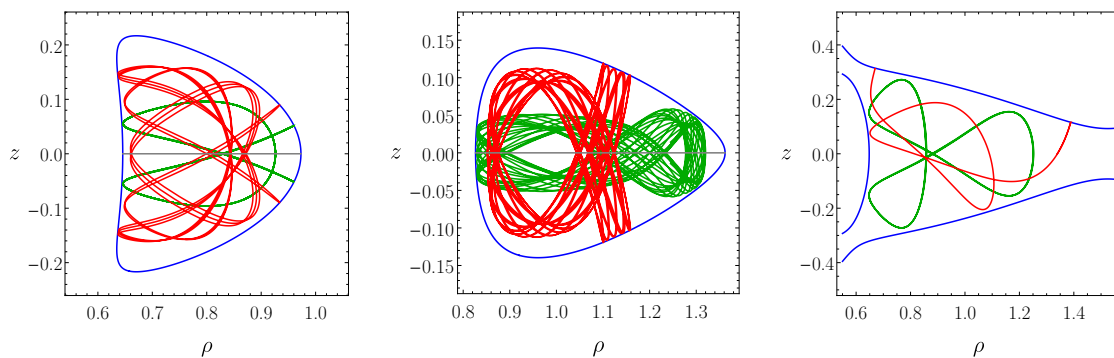


Figure 6.5. (a) Poincaré section in the (ρ, p_ρ) -plane for $\mu = 0.4$. (b) A pair of bounded orbits for $\mu = 0.4$: a low-order resonant orbit [red]; a chaotic orbit [green]. (c) Poincaré section in the (ρ, p_ρ) -plane for $\mu = 0$, with chaos dominant. There exist small “islands of stability” in the “chaotic sea”. (d) A pair of bounded orbits for $\mu = 0$: a $2 : 2$ resonant orbit [red]; a chaotic orbit [green].


 (a) Poincaré section for $d = 0.95M$, $\mu = 0.2$

 (b) Poincaré section for $d = 1.05M$, $\mu = 0.05$

Figure 6.6. Poincaré sections for the “anisotropic” ($d \neq M$) equal-mass Majumdar–Papapetrou di-hole. (a) An example with $d < M$. The Poincaré section exhibits rich structure: elliptic and hyperbolic fixed points; low-order resonances; KAM islands; chaotic bands; and novel features, including “crenulations”. (b) An example with $d > M$ which exhibits regular behaviour.


 (a) $d = 0.95M$, $\mu = 0.20$

 (b) $d = 1.05M$, $\mu = 0.05$

 (c) $d = M$, $\mu = -0.10$

Figure 6.7. (a) Regular bounded null orbits for the $d < M$ case, associated with the “crenulations” in the Poincaré section of Figure 6.6(a). (b) Regular bounded null orbits for the $d > M$ case; see Figure 6.6(b) for the corresponding Poincaré section. (c) Non-escaping orbits in the unbounded regime with $\mu = -0.1$: an open libration-type orbit which touches the contour $h = p_\phi$ [red]; a closed rotation-type orbit [green]. These orbits persist, despite the existence of three narrow escapes which connect the scattering region to the black holes and to infinity (cf. Chapter 5).

are manifest around resonances [103]. Figure 6.5(b) shows a 5 : 5 resonant orbit [red]; and a chaotic orbit [green]. As μ is decreased, further lower-order resonances bifurcate from the 1 : 1 resonances exhibited in Figure 6.4(b); and we observe the transition from order to chaos. At the threshold $\mu = 0$, chaos is dominant – as shown by the Poincaré section in Figure 6.5(c). We see that almost all of the Poincaré section is filled by chaotic orbits, with small “islands of stability” manifest. Figure 6.5(d) shows a pair of orbits bounded by the contour $h = p_\phi^*$. In particular, we show a 2 : 2 resonant orbit [red], and an example of a chaotic orbit [green].

Intriguingly, the qualitative features of kinematically bounded photon orbits around the isotropic ($d = M$) equal-mass Majumdar–Papapetrou di-hole are shared by the well-known Hénon–Heiles Hamiltonian system [103, 117, 296], a review of which is given in Appendix C. In particular, compare the Poincaré sections and bounded orbits in Figures 6.4 and 6.5 (Majumdar–Papapetrou) with those of Figures C.4 and C.5 (Hénon–Heiles).

Figure 6.6 shows Poincaré sections for the “anisotropic” ($d \neq M$) equal-mass Majumdar–Papapetrou di-hole. We observe a rich variety of features in the $d < M$ case, shown in Figure 6.6(a). In addition to elliptic and hyperbolic fixed points (periodic orbits), there exist chaotic domains in which KAM islands of stability are embedded. Moreover, we observe novel features in this Poincaré section, including the “crenulations”. The Poincaré section shown in Figure 6.6(b) ($d > M$) exhibits regular behaviour, despite the fact that μ is small (i.e., we are far from the maximum of h).

Intriguingly, non-escaping photon orbits survive even in the kinematically unbounded regime where $\mu < 0$ (i.e., $p_\phi < p_\phi^{(1)}$). Figure 6.7(c) shows examples of closed (rotation-type) and open (libration-type) periodic orbits in for $\mu = -0.1$. Recall from Section 5.2.2 that, for $\mu \lesssim 0$ there exist KAM tori of quasi-periodic orbits that do not escape the scattering region, despite the system being *open* (i.e., kinematically unbounded); see the exit basins of Figures 5.2(c) and 5.2(f), for example.

6.6 Discussion

6.6.1 Extensions

Stable photon orbits in other spacetimes

Cunha *et al.* [184] explore null geodesic motion around boson stars and Kerr black holes with scalar hair. Two-dimensional effective potentials – similar to those discussed in Section 6.3.1 – are employed, since the motion is not known (nor expected) to be separable. The authors find that the spacetimes admit stable light-rings, which permit the formation of “pocket” features in the two-dimensional configuration space spanned by the Boyer–Lindquist coordinates $\{r, \theta\}$; these pocket features are delimited by the contours of the effective potential. As in the Majumdar–Papapetrou di-hole case, such features lead to an effective trapping of some null rays. Furthermore, the authors observe that the existence of multiple light-rings (including a stable one) around a horizon is a central ingredient for the existence of multiple disconnected shadows of a single hairy black hole.

Using techniques from the theory of dynamical systems, Grover and Wittig [297] show

how light-rings – fixed points of the null geodesic potential – give rise to families of periodic orbits and invariant manifolds in phase space. These invariant structures are shown to define the shape of the black hole shadow as well as a number of key signatures of gravitational lensing. As a case study, the authors analyse chaotic lensing by a Herdeiro–Radu hairy black hole (i.e., a Kerr black hole with scalar hair), similar to those studied by Cunha *et al.* [191, 241].

In a recent study, Jia *et al.* [298] explore the existence and stability of null (and timelike) equatorial circular orbits in *static* axisymmetric spacetimes. Using a fixed point approach similar to that considered here, the authors obtain a necessary and sufficient conditions for the (non-)existence of timelike circular orbits; they prove that (stable) timelike circular orbits will always exist at large ρ in an asymptotically flat static axisymmetric spacetime with a positive Arnowitt–Deser–Misner mass. The authors then derive the necessary and sufficient condition for the existence of null circular orbits. They show that the existence of timelike circular orbits in static axisymmetric spacetimes does not, in general, imply the existence of null circular orbits, and vice versa.

To generalise the results for stationary axisymmetric electrovacuum spacetimes (Section 6.3.2), one could consider (i) including a non-zero cosmological constant; (ii) other fields and matter sources; or (iii) spacetimes with less symmetry. Let us comment on each of these potential generalisations here.

A natural starting point for case (i) would be to consider a cosmological version of the Weyl–Lewis–Papapetrou solution: the spacetime geometry and an Ernst-type formulation of the field equations are presented in vacuum by Charmousis *et al.* [299], and in electrovacuum by Astorino [300]. In both cases, the inclusion of a non-zero cosmological constant modifies the metric and field equations in a non-trivial way; a careful generalisation of the results of Section 6.3.2 would therefore be necessary in this case. As a first attempt at case (ii), one may wish to analyse the existence and stability of photon orbits in stationary axisymmetric (electro)vacuum spacetimes which are coupled to a scalar field. Solutions in the case of conformally (and minimally) coupled scalar fields in vacuum are considered by Astorino [301]. We caution that, even in the “simple” case of massless minimally or conformally couple scalar fields, the metric and field equations are much less tractable than in the case of stationary axisymmetric electrovacuum with no additional matter fields (considered in this work). Finally, we remark that, whilst case (iii) is of interest for dynamical binary black holes (which are expected to be neither stationary nor axisymmetric), the lack of symmetry would mean that studying the existence and stability of photon orbits using two-dimensional effective potentials would not be possible. A special treatment would be required to study the existence and stability of circular photon orbits in dynamical spacetimes. To gain insight into the nature of photon orbits on the spacetimes of dynamical binary black holes, a systematic study of the simulations generated by Bohn *et al.* [190], for example, could repay further investigation.

Light-ring stability for horizonless ultra-compact objects

In Section 6.3.2, we considered the stability of photon orbits for (electro)vacuum solutions to the Einstein(–Maxwell) field equations, motivated by the study of black holes. In

recent work, Cunha *et al.* [268] consider stationary axisymmetric solutions of Einstein’s field equations that are formed from classical gravitational collapse of matter which obeys the null energy condition, which are ultra-compact (i.e., they possess a light-ring), but which do not possess an event horizon.

Using a Hamiltonian formalism for null geodesics similar to that of Sections 6.3.1 and 6.3.2, the authors derive a pair of effective potentials which they use to analyse the existence and stability of circular photon orbits. The authors prove that, for ultra-compact objects described by a stationary axisymmetric geometry which is a solution to *any* metric theory of gravity in which photons propagate along null geodesics, light-rings always come in pairs, one of which is a saddle point and the other is a local extremum of an effective potential. The proof depends on a topological argument based on the Brouwer degree of a continuous map. Then, assuming Einstein’s equations of general relativity, the authors prove that the local extremum of the effective potential is a *minimum* (i.e., a stable light-ring) if the stress–energy tensor satisfies the null energy condition.

In a follow-up paper, Hod [302] proves that, whilst the key theorem of Cunha *et al.* [268] is generally true, there is an exception in the case of *degenerate light-rings*, which are characterised by the simple relation $8\pi(\rho + p_T)r_\gamma^2 = 1$, where $r = r_\gamma$ is the location of the light-ring in spherical coordinates, $\rho = -T^t_t$ is the energy density, and $p_T = T^\theta_\theta = T^\phi_\phi$ is the tangential pressure. This is illustrated explicitly by considering the case of a spherically symmetric constant-density stars with dimensionless compactness $\frac{M}{R} = \frac{1}{3}$, where M is the object’s mass and R is its surface radius; Hod finds that this system admits a unique *unstable* light-ring.

6.6.2 Conclusions

In this chapter, we have discussed the existence, stability and phenomenology of circular photon orbits (or light-rings) in the context of four-dimensional stationary axisymmetric solutions to the Einstein–Maxwell equations of gravity and electromagnetism.

Geodesic motion on Kerr(–Newman) spacetime, which is perhaps the most astrophysically relevant stationary axisymmetric (electro)vacuum solution, is separable (and hence Liouville integrable) in Boyer–Lindquist coordinates, thanks to the existence of four integrals of motion in involution on phase space. The fourth conserved quantity is the Carter constant, which arises from the “hidden” symmetry of Kerr(–Newman) spacetime which is encoded in the existence of a rank-two Killing tensor field (see Section 2.6.4). The separability of the radial and latitudinal motion in Boyer–Lindquist coordinates means that the problem of classifying the equatorial circular photon orbits of Kerr–Newman spacetime reduces to classifying the roots of the quartic (6.7) in $u = \frac{M}{r}$, where r is the radial coordinate and M is the mass. We reviewed the Balek–Bičák–Stuchlík phase diagram [78] in the charge–spin plane for equatorial circular photon orbits of positive Boyer–Lindquist radius r (Figure 6.1).

We focussed on stationary axisymmetric spacetimes in four dimensions, which are described in Weyl–Lewis–Papapetrou coordinates $\{t, \rho, z, \phi\}$ by the line element (6.11). In general, the geodesic motion on such geometries is not expected to be separable (except in special cases, such as the Kerr–Newman spacetime); however, the stationarity and axisym-

metry mean that t and ϕ are ignorable coordinates, and the motion of null rays is governed by a two-dimensional time-independent Hamiltonian system with conserved parameter p_ϕ . We introduced the effective potentials $h^\pm(\rho, z)$, given by (6.17), and demonstrated that the problem of classifying circular photon orbits reduces to that of classifying the stationary points of these effective potentials.

Restricting our attention to the electrovacuum case, we classified the stationary points of the effective potentials by availing a subset of the Einstein–Maxwell field equations. Our key result is that, in the context of four-dimensional stationary axisymmetric spacetimes in general relativity, generic stable photon orbits are forbidden in pure vacuum; however, stable photon orbits may arise in *electrovacuum*, thanks to (6.32).

As a case study, we explored the existence of stable photon orbits around static di-holes in four dimensions. We reviewed the general Bretón–Manko–Aguilar family of solutions (Section 6.4.1), focussing on the Reissner–Nordström di-hole subfamily, which contains the Weyl–Bach and Majumdar–Papapetrou di-holes as special cases. We found that, in the case of the highly symmetric equal-mass Majumdar–Papapetrou di-hole, stable photon orbits exist for separations in the range $\sqrt{\frac{16}{27}} < \frac{d}{M} < \sqrt{\frac{32}{27}}$. In cases with less symmetry (e.g. the unequal-mass Majumdar–Papapetrou di-hole and the equal-mass, equal-charge Reissner–Nordström di-hole), we used a numerical method to determine the existence region for stable photon orbits in parameter space (Figure 6.3). The results of Figure 6.3(b) indicate that stable photon orbits can exist up to (but not including) the vacuum Weyl–Bach case, provided that other parameters are appropriately fine-tuned.

Finally, we explored the rich structure of stable photon orbits for the Majumdar–Papapetrou di-hole, which is non-integrable (Section 6.5). In the isotropic case with separation $d = M$, we found that the system can either be closed or open with three narrow escapes, depending on the value of the conserved azimuthal angular momentum. Close to the minimum of the geodesic potential, the motion is bounded to a small elliptical region. Using Poincaré sections, we analysed the dynamics of bounded null geodesics, presented some examples of resonant orbits, and commented on the transition from order to chaos as the azimuthal angular momentum parameter is varied. Curiously, in both the bounded regime (this chapter) and the unbounded regime (Chapter 5), the isotropic equal-mass Majumdar–Papapetrou di-hole shares many qualitative features with the Hénon–Heiles Hamiltonian system (Appendix C). In the anisotropic case ($d \neq M$), we observe a variety of rich structure in the Poincaré sections (e.g. the “crenulations” of Figure 6.6) and the bounded null geodesics (Figure 6.7).

Chapter 7

Higher-order geometric optics on Kerr spacetime

7.1 Introduction

The Event Horizon Telescope’s observations of the black hole M87* [1] and the detections of gravitational waves by LIGO–Virgo [18] have opened up an exciting new era in gravitational astronomy (see Chapter 1). These events are part of a long line of observations which involve measurements of (typically electromagnetic) waves that have travelled to us over cosmological distances through a dynamical curved spacetime.

Formally, electromagnetic fields are described by wavelike solutions to Maxwell’s equations; however, these equations are not typically analysed directly in practice. To describe the effects of gravitational lensing, it is usually sufficient to adopt a leading-order geometric optics approximation [7, 29, 162]. In this approach, the gradient of the eikonal phase is tangent to a light ray, and light travels along null geodesics in a four-dimensional spacetime, which is typically taken to be a solution to Einstein’s field equations of general relativity; variations in the intensity of the radiation are determined by the change in cross-sectional area of an infinitesimal bundle of rays in the geometric optics wavefront; and the polarisation of the wave is parallel-transported along the ray – the gravitational Faraday effect [303].

Geometric optics does not only apply to electromagnetic waves and the description of gravitational lensing phenomena; in fact, the formalism is widely used in many areas of physics [304]. Geometric optics is an accurate description of wave propagation when the wavelengths and inverse frequency scales are short when compared to other characteristic length and time scales, including the spacetime curvature scale set by the Riemann tensor [29]. In the context of general relativity, the leading-order geometrical optics formalism has been employed to provide a theoretical description of scalar, electromagnetic and gravitational fields on curved spacetime [171].

For observations of gravitational lensing effects, the assumptions on which the geometric optics approximation relies are generally sound. For example, the Event Horizon Telescope employs millimetre and sub-millimetre very-long-baseline interferometry to generate images of supermassive black holes with very large diameters ($\sim 10^{13}$ m for M87*).

Despite the fact that gravitational waves have much longer wavelengths ($\sim 10^7$ m for GW150914), the geometric optics approximation is still typically safe.

However, there are a number of situations where the geometric optics approximation breaks down. For example, it is well-known that – even in standard optics – the geometric optics approximation is not valid at caustics (i.e., transverse self-intersections of the wavefront) [7]. There are other scenarios, such as the scattering of waves by Kerr black holes when the wavelength is comparable with the horizon radius [305, 306], where the leading-order geometric optics formalism is no longer an accurate description of the wave-optical phenomena. In the latter scenario, one would expect that higher-order corrections would provide a more accurate description, and perhaps provide insight into wavelike effects which are not offered by the leading-order (ray-optics) approximation.

Higher-order extensions to the geometric optics formalism have been studied by Ehlers [169] and Anile [170]. More recently, modifications to the geometric optics formalism have been given by Dolan [81] and Harte [307, 308]. The work of Dolan [81] is motivated by interest in spin–helicity effects, i.e., a coupling between the helicity of a circularly polarised electromagnetic wave (of finite wavelength) and the frame-dragging of spacetime outside a rotating source, such as a Kerr black hole [309]. It has been demonstrated, for example, that Kerr black holes are able to distinguish and separate waves of opposite helicity [305, 310, 311]. Moreover, recent analyses of the absorption of planar electromagnetic waves incident on a Kerr black hole along the axis of symmetry demonstrated that waves with a circular polarisation that is rotating in the opposite sense to the black hole are absorbed to a greater degree than co-rotating waves [306]. Dolan [81] argued that, in principle, such effects can be captured by an extension to geometric optics at sub-leading order. However, this approach has not yet been employed in any practical calculations of lensing effects by black holes.

In this chapter, we apply the higher-order geometric optics formalism to electromagnetic waves on Kerr spacetime. In Section 7.2 we review the geometric optics formalism for electromagnetic waves on a four-dimensional curved spacetime. The leading-order formalism was presented in detail in Section 3.3.2, but is reviewed in Section 7.2.1 of the present chapter for convenience. In Section 7.2.2, we summarise the higher-order geometric optics formalism of Dolan [81]. In this work, we are interested in using the formalism to calculate the sub-leading-order corrections to the electromagnetic stress–energy tensor for circularly polarised waves propagating on Kerr spacetime. We review the important properties of Kerr spacetime in Section 7.3, including the spacetime geometry, the existence of Killing objects and symmetries, Carter’s symmetric tetrad, and the role of spacetime symmetries in the separability of geodesic motion. As described in Section 3.3.2, (parallel-transported) complex null tetrads play an important role in the analysis of gravitational lensing phenomena through the geometric optics formalism. In Section 7.4, we use the principal tensor and its Hodge dual to construct a complex null tetrad which is parallel-propagated along the null geodesics of Kerr spacetime; we demonstrate that this tetrad is equivalent to a complex null tetrad which can be built from the legs of a real parallel-propagated tetrad found by Marck [52]. Section 7.5 concerns the Newman–Penrose formalism for our parallel-propagated complex null tetrad: we present the transport equations required to compute

the higher-order geometric optics corrections; and we calculate quantities such as the Weyl curvature scalars and their directional derivatives along the legs of the complex null tetrad. In Section 7.6, we analyse the far-field behaviour of the Weyl scalars, Newman–Penrose scalars, and geometric optics quantities. In Section 7.7, we review wavefronts and caustics in geometric optics, and analyse the behaviour of (divergent) Newman–Penrose quantities in the neighbourhood of caustic points, where neighbouring rays cross. Finally, in Section 7.8, we present a practical method for evolving transport equations through caustic points: we introduce a regularisation method in Section 7.8.1; and we present some preliminary numerical results for Kerr spacetime in Section 7.8.2. We discuss our results and avenues for future work in Section 7.9.

7.2 Higher-order geometric optics formalism

7.2.1 Leading-order geometric optics

In this section, we briefly review aspects of the leading-order geometric optics formalism. For further details, see Section 3.3.2 and references therein. The Faraday tensor F_{ab} is governed by the Maxwell equations in curved spacetime. One may introduce a complexified Faraday tensor $\mathcal{F}_{ab} = F_{ab} + i^*F_{ab}$, which is a self-dual bivector. Here, $^*F_{ab}$ is the Hodge dual of the Faraday tensor, defined in (3.22). In the absence of sources, the complexified Faraday tensor satisfies the wave equation

$$\square\mathcal{F}_{ab} + 2R_{abcd}\mathcal{F}^{cd} + R_a{}^c\mathcal{F}_{bc} - R_b{}^c\mathcal{F}_{ac} = 0, \quad (7.1)$$

which is obtained by taking further covariant derivatives of the source-free Maxwell's equations (3.26).

The leading-order geometric optics solution to (7.1) for a circularly polarised electromagnetic wave is given by

$$\mathcal{F}_{ab} = \mathcal{A}f_{ab} \exp(i\omega\Phi), \quad f_{ab} = 2k_{[a}m_{b]}, \quad (7.2)$$

Here, k^a is a real null ($k^ak_a = 0$) vector field which is defined as the gradient of the eikonal phase ($k_a = \nabla_a\Phi$). We denote the directional derivative along k^a by $D = k^a\nabla_a$. From the gradient property, it follows that $k_{[a;b]} = 0$. The fact that k^a is null and a gradient means k^a is tangent to a geodesic ($Dk^a = 0$); k^a is a null generator of constant-phase hypersurfaces ($\Phi = \text{constant}$). The vector m^a and its complex conjugate \bar{m}^a are complex null ($m^am_a = 0$) vector fields which are unit $m^a\bar{m}_a = 1$; parallel-propagated ($Dm^a = 0$); transverse to k^a ($k^am_a = 0$); and hence tangent to constant-phase hypersurfaces ($m^a\nabla_a\Phi = 0$). One may introduce an auxiliary real null vector n^a which is future-pointing and satisfies $k^an_a = -1$ and $m^an_a = 0$. It follows from these scalar products that, if k^a and m^a are parallel-transported, then so too is n^a ($Dn^a = 0$). Hence, each leg of the complex null tetrad $\{k^a, n^a, m^a, \bar{m}^a\}$ is parallel-propagated along the null generators of the constant-phase hypersurfaces.

We recall from Section 3.3.2 that, at $O(\omega)$, the ambiguity in the definitions of the amplitude \mathcal{A} and the polarisation bivector f_{ab} can be used to split the transport equations

such that

$$D\mathcal{A}^2 = 2\rho\mathcal{A}^2, \quad D\mathfrak{f}_{ab} = 0, \quad (7.3)$$

where $\rho = -m^a k_{a;b} \bar{m}^b$ is a Newman–Penrose scalar (see Section 2.5.3). Hereafter, we will use the transport equation for the square-amplitude \mathcal{A}^2 rather than \mathcal{A} for reasons which shall be discussed in Section 7.7.

At leading order in ω , the electromagnetic stress–energy tensor takes the form of a null fluid, i.e.,

$$T_{ab} = \frac{1}{8\pi} \mathcal{A}^2 k_a k_b. \quad (7.4)$$

We recall that, for a twist-free null congruence ($k_{[a;b]} = 0$), the cross-sectional area A of an infinitesimal bundle of rays in the geometric-optics wavefront satisfies the transport equation

$$DA = -2\rho A, \quad (7.5)$$

where the Newman–Penrose scalar (or optical scalar) ρ , which is responsible for the expansion of the infinitesimal bundle, is real.

7.2.2 Higher-order geometric optics

Dolan [81] presents a method to extend the geometrical approximation to higher orders in the frequency ω . This is achieved by retaining the standard ansatz (3.30), and expanding the polarisation bivector as a power series:

$$\mathfrak{f}_{ab} = \sum_{j=0}^{\infty} \omega^{-j} \mathfrak{f}^{(j)}_{ab}. \quad (7.6)$$

The self-dual bivectors $\mathfrak{f}^{(j)}_{ab}$ are then expanded in the bivector basis $U_{ab} = 2k_{[a} m_{b]}$, $V_{ab} = 2\bar{m}_{[a} n_{b]}$, $W_{ab} = 2(m_{[a} \bar{m}_{b]} - k_{[a} n_{b]})$, which is constructed from a twist-free parallel-propagated complex null tetrad $\{k^a, n^a, m^a, \bar{m}^a\}$ [83], viz.

$$\mathfrak{f}^{(j)}_{ab} = \mathbf{u}_j U_{ab} + \mathbf{v}_j V_{ab} + \mathbf{w}_j W_{ab}, \quad (7.7)$$

where \mathbf{u}_j , \mathbf{v}_j and \mathbf{w}_j are complex scalar fields.

Inserting the ansatz (7.6) into the first of Maxwell’s equations (3.26) and expanding order-by-order in ω yields a system of equations for the $\mathfrak{f}^{(j)}_{ab}$. At sub-leading order, the relevant equations are

$$D\mathbf{u}_1 = i \left(\frac{1}{\mathcal{A}} \bar{\delta} \delta \mathcal{A} + \frac{\chi}{\mathcal{A}} \bar{\delta} \mathcal{A} + \bar{\delta} \chi - \sigma \lambda \right), \quad (7.8)$$

$$\mathbf{v}_1 = i\sigma, \quad (7.9)$$

$$\mathbf{w}_1 = i \left(\frac{1}{\mathcal{A}} \delta \mathcal{A} + \chi \right), \quad (7.10)$$

where $\sigma = -m^a k_{a;b} m^b$ and $\chi = \bar{m}^a m_{a;b} m^b$ are Newman–Penrose scalars (see Sections 2.5 and 7.5). The transport equation (7.8) features second derivatives of the amplitude \mathcal{A} across the wavefront (i.e., in the m - and \bar{m} -directions).

Using the expression for the electromagnetic stress–energy tensor (3.29), and identities for the bivector basis $\{U_{ab}, V_{ab}, W_{ab}\}$, the sub-leading-order correction to the stress–energy tensor is

$$4\pi T_{ab} = \frac{1}{2}\mathcal{A}^2 k_a k_b + \omega^{-1}\mathcal{A}^2 \operatorname{Re}\left(\mathbf{u}_1 k_a k_b + \mathbf{v}_1 \bar{m}_a \bar{m}_b - 2\mathbf{w}_1 k_{(a} \bar{m}_{b)}\right) + O(\omega^{-2}). \quad (7.11)$$

Since this only features the real part of \mathbf{u}_1 , one may use (7.8) and its complex conjugate to obtain a transport equation for $\operatorname{Re}(\mathbf{u}_1)$. One may then apply the commutator identity $\bar{\delta}\delta - \delta\bar{\delta} = (\bar{\mu} - \mu)D - \bar{\chi}\delta + \chi\bar{\delta}$ to eliminate second derivatives of \mathcal{A} in favour of first derivatives. The relevant transport equation is then

$$D(\operatorname{Re}(\mathbf{u}_1)) = \frac{i}{2}\left[\bar{\delta}\chi - \delta\bar{\chi} + \chi\frac{\bar{\delta}\mathcal{A}^2}{\mathcal{A}^2} - \bar{\chi}\frac{\delta\mathcal{A}^2}{\mathcal{A}^2} + \rho(\bar{\mu} - \mu) + \bar{\sigma}\bar{\lambda} - \sigma\lambda\right]. \quad (7.12)$$

Our principal aim here is to calculate the sub-dominant correction to the stress–energy tensor (7.11). In principle, the Newman–Penrose scalars which appear in (7.9), (7.10) and (7.12) may be found along null geodesics by evolving a system of standard transport equations [81, 83] (see Sections 2.5.5 and 7.5.1). We observe that (7.9), (7.10) and (7.12) also feature “higher-order” Newman–Penrose quantities and geometric optics quantities, e.g. $\bar{\delta}\chi$ and $\delta\mathcal{A}^2$, for which transport equations may also be obtained (see Section 7.5.1 and [81]).

By considering the asymptotic behaviour of the Newman–Penrose quantities in a flat region of spacetime, Dolan [81] finds that, as $s \rightarrow \infty$, $\mathbf{u}_j = O(1)$, $\mathbf{v}_j = O(s^{-2})$ and $\mathbf{w}_j = O(s^{-1})$, where s denotes the affine parameter along a null geodesic. To find the $O(\omega^{-1})$ correction to T_{ab} in (7.11) far from the black hole, we therefore only need to calculate the quantity $\operatorname{Re}(\mathbf{u}_1)$ by means of (7.12). We note that the sub-leading-order correction is proportional to the leading-order piece (i.e., $k_a k_b$).

7.3 Kerr spacetime

7.3.1 Spacetime geometry

The Kerr solution to Einstein’s field equations describes the spacetime of a rotating black hole [35]. The region outside the event horizon may be described in Boyer–Lindquist coordinates $\{t, r, \theta, \phi\}$ [46], in which the line element $d\bar{s}^2 = g_{ab}dx^a dx^b$ takes the form¹

$$\begin{aligned} d\bar{s}^2 = & -\left(1 - \frac{2Mr}{\Sigma}\right)dt^2 + \frac{\Sigma}{\Delta}dr^2 + \Sigma d\theta^2 + \left(r^2 + a^2 + \frac{2Ma^2r}{\Sigma}\sin^2\theta\right)\sin^2\theta d\phi^2 \\ & - \frac{4Mar\sin^2\theta}{\Sigma}dt d\phi, \end{aligned} \quad (7.13)$$

where $\Sigma(r, \theta) = r^2 + a^2 \cos^2\theta$ and $\Delta(r) = r^2 - 2Mr + a^2$. Here, M is the mass of the black hole and $a = \frac{J}{M}$ is the spin, where J is the angular momentum. The Kerr spacetime has an outer (+) and inner (−) horizon at $r_{\pm} = M \pm \sqrt{M^2 - a^2 \cos^2\theta}$; we only consider the exterior region $r > r_+$. The non-zero components of the inverse metric g^{ab} are

$$g^{tt} = -\frac{1}{\Delta}\left(r^2 + a^2 + \frac{2Mr a^2}{\Sigma}\sin^2\theta\right), \quad g^{t\phi} = g^{\phi t} = -\frac{2Mar}{\Delta\Sigma}, \quad (7.14)$$

¹Here, we denote the spacetime interval by $d\bar{s}$, rather than ds , to avoid a clash with the notation used for the affine parameter throughout the rest of this chapter.

$$g^{rr} = \frac{\Delta}{\Sigma}, \quad g^{\theta\theta} = \frac{1}{\Sigma}, \quad g^{\phi\phi} = \frac{\Delta - a^2 \sin^2 \theta}{\Delta \Sigma \sin^2 \theta}. \quad (7.15)$$

It is useful to note the identity $\Delta - a^2 \sin^2 \theta = \Sigma - 2Mr$.

7.3.2 Killing objects and symmetries

In Boyer–Lindquist coordinates, the Kerr solution admits a pair of Killing vectors $\xi_{(t)} = \partial_t$ and $\xi_{(\phi)} = \partial_\phi$, which satisfy Killing’s equation,

$$\mathcal{L}_\xi g_{ab} = \xi_{a;b} + \xi_{b;a} = 2\xi_{(a;b)} = 0. \quad (7.16)$$

It is well known that the Killing vectors $\xi_{(t)}$ and $\xi_{(\phi)}$ encode continuous symmetries (isometries) of the metric, corresponding to time-translations and axial rotations, respectively.

In fact, the Kerr metric admits a special geometric object – the *principal tensor* – which is the generator of a complete set of explicit and hidden (or implicit) symmetries of the spacetime [50]. As described in Section 2.6.4, the principal tensor h_{ab} is a closed ($dh = 0$) conformal Killing–Yano two-form, which satisfies the defining equation

$$h_{ab;c} = g_{ac}\xi_b - g_{cb}\xi_a, \quad (7.17)$$

where ξ^a is an associated vector field. (The terms “principal tensor” and “closed conformal Killing–Yano tensor” will be used interchangeably.) Contracting (7.17) with g^{cb} and using the antisymmetry of h_{ab} , one may show that

$$\xi_a = \frac{1}{3}h_{ba}{}^{;b}. \quad (7.18)$$

Taking the covariant derivative of (7.18), symmetrising, and imposing the Einstein field equations in four dimensions, it is possible to show that [106, 312, 313]

$$\xi_{(a;b)} = \frac{1}{2}R_{c(b}h_a)^c. \quad (7.19)$$

For a spacetime which satisfies Einstein’s field equations in vacuum with a cosmological constant, the Ricci tensor is proportional to the metric tensor. Hence, due to the antisymmetry of h_{ab} , the right-hand side of (7.19) vanishes, and we see that $\xi_{(a;b)}$ i.e., $\mathcal{L}_\xi g_{ab} = 0$. In other words, ξ^a defined through (7.18) is a Killing vector. In fact, any solution to Einstein’s equations which admits a principal tensor h_{ab} automatically possesses a Killing vector. This is referred to as the *principal Killing vector*. For Kerr, the principal Killing vector is $\xi_{(t)} = \partial_t$ in Boyer–Lindquist coordinates.

The Hodge dual of the closed conformal Killing–Yano tensor is the *Killing–Yano tensor* $f_{ab} = {}^*h_{ab} = \frac{1}{2}\epsilon_{abcd}h^{cd}$, which satisfies $f_{ab;c} = f_{[ab;c]}$. The *Killing tensor* $K_{ab} = f_a{}^c f_{bc}$ is the “square” of the Killing–Yano tensor, and satisfies $K_{(ab;c)} = 0$. The *conformal Killing tensor* $Q_{ab} = h_a{}^c h_{bc}$ is the “square” of the closed conformal Killing–Yano tensor. This tensor satisfies $Q_{(ab;c)} = g_{(ab}q_c)$, where $q_a = h_{ab}\xi_{(t)}{}^b$.

In this chapter, we focus on the null geodesics of Kerr spacetime. Consider an affinely parametrised null geodesic on Kerr spacetime with tangent vector field $k^a(s) = \dot{x}^a(s)$, where an overdot denotes differentiation with respect to the affine parameter s . In Section 2.6.4, we saw that each Killing vector ξ^a yields a conserved quantity along a geodesic,

$I = g_{ab}\xi^a k^b$. On Kerr spacetime, there exists a pair of commuting Killing vectors, $\xi_{(t)}$ and $\xi_{(\phi)}$, which encode the stationarity and axisymmetry of the metric. These Killing vectors yield a pair of conserved quantities, $E = -g_{ab}\xi_{(t)}^a k^b$ and $L_z = g_{ab}\xi_{(\phi)}^a k^b$, such that $\dot{E} = 0 = \dot{L}_z$ along rays.

Besides the conserved quantities which arise from Killing vectors, there also exist conserved quantities which are related to higher-rank Killing tensors, which imply the existence of less evident symmetries [50]. For Kerr spacetime, define the scalar quantities

$$K = K_{ab}k^a k^b, \quad Q = Q_{ab}k^a k^b, \quad (7.20)$$

along a null geodesic with tangent vector field k^a . Then

$$\dot{K} = k^c K_{;c} = K_{ab;c} k^a k^b k^c = K_{(ab;c)} k^a k^b k^c = 0, \quad (7.21)$$

where the right-hand side vanishes due to the fact that $K_{(ab;c)} = 0$, i.e., K_{ab} is a Killing tensor. Similarly,

$$\dot{Q} = k^c Q_{;c} = Q_{ab;c} k^a k^b k^c = Q_{(ab;c)} k^a k^b k^c = g_{(ab} q_c) k^a k^b k^c = k^a k_a k^c q_c = 0, \quad (7.22)$$

by virtue of the fact that k^a is null. Hence, the Killing tensor and conformal Killing tensor yield conserved quantities (7.20) along null geodesics. These will be discussed in greater detail in Section 7.3.4.

7.3.3 Carter's canonical tetrad

Here, we introduce Carter's "canonical" (symmetric) tetrad [47], an orthonormal tetrad (see Section 2.5.1) which is useful when analysing the separability properties of Kerr spacetime [291, 314]. The one-form components of the tetrad $\omega^{(\alpha)}_a$ are given by

$$\omega^{(0)}_a dx^a = \sqrt{\frac{\Delta}{\Sigma}} (dt - a \sin^2 \theta d\phi), \quad (7.23)$$

$$\omega^{(1)}_a dx^a = \sqrt{\frac{\Sigma}{\Delta}} dr, \quad (7.24)$$

$$\omega^{(2)}_a dx^a = \sqrt{\Sigma} d\theta, \quad (7.25)$$

$$\omega^{(3)}_a dx^a = \frac{\sin \theta}{\sqrt{\Sigma}} [a dt - (r^2 + a^2) d\phi], \quad (7.26)$$

and the vector components of the tetrad $\omega_{(\alpha)}^a$ are

$$\omega_{(0)}^a \partial_a = \frac{1}{\sqrt{\Sigma \Delta}} [(r^2 + a^2) \partial_t + a \partial_\phi], \quad (7.27)$$

$$\omega_{(1)}^a \partial_a = \sqrt{\frac{\Delta}{\Sigma}} \partial_r, \quad (7.28)$$

$$\omega_{(2)}^a \partial_a = \frac{1}{\sqrt{\Sigma}} \partial_\theta, \quad (7.29)$$

$$\omega_{(3)}^a \partial_a = -\frac{1}{\sqrt{\Sigma} \sin \theta} (a \sin^2 \theta \partial_t + \partial_\phi). \quad (7.30)$$

In terms of Carter's orthonormal tetrad, the covariant and contravariant components of the metric tensor are, respectively, given by

$$g_{ab} = -\omega^{(0)}{}_a \omega^{(0)}{}_b + \omega^{(1)}{}_a \omega^{(1)}{}_b + \omega^{(2)}{}_a \omega^{(2)}{}_b + \omega^{(3)}{}_a \omega^{(3)}{}_b, \quad (7.31)$$

$$g^{ab} = -\omega_{(0)}{}^a \omega_{(0)}{}^b + \omega_{(1)}{}^a \omega_{(1)}{}^b + \omega_{(2)}{}^a \omega_{(2)}{}^b + \omega_{(3)}{}^a \omega_{(3)}{}^b. \quad (7.32)$$

Carter's tetrad is closely related to the (normalised) *Darboux basis* $\{\nu, \hat{\nu}, \epsilon, \hat{\epsilon}\}$. This is presented by Frolov *et al.* [50] in *canonical coordinates* $\{\tau = t - a\phi, r, y = a \cos \theta, \psi = \frac{\phi}{a}\}$, in which the hidden symmetry of Kerr spacetime is more evident. By direct comparison of the Darboux basis in [50] with the basis one-forms of Carter's symmetric tetrad (7.23)–(7.26), we see that

$$\nu = \omega^{(1)}, \quad \hat{\nu} = \omega^{(0)}, \quad \epsilon = -\omega^{(2)}, \quad \hat{\epsilon} = \omega^{(3)}. \quad (7.33)$$

Hence, Carter's orthonormal tetrad is equivalent to the Darboux frame (up to a change of sign). One advantage of the Darboux basis (i.e., Carter's tetrad) is that the closed conformal Killing–Yano tensor and related Killing objects take a remarkably simple form [50]:

$$f_{ab} = -2a \cos \theta \omega^{(0)}{}_{[a} \omega^{(1)}{}_{b]} - 2r \omega^{(2)}{}_{[a} \omega^{(3)}{}_{b]}, \quad (7.34)$$

$$h_{ab} = 2r \omega^{(0)}{}_{[a} \omega^{(1)}{}_{b]} - 2a \cos \theta \omega^{(2)}{}_{[a} \omega^{(3)}{}_{b]}, \quad (7.35)$$

$$K_{ab} = -a^2 \cos^2 \theta \left(-\omega^{(0)}{}_a \omega^{(0)}{}_b + \omega^{(1)}{}_a \omega^{(1)}{}_b \right) + r^2 \left(\omega^{(2)}{}_a \omega^{(2)}{}_b + \omega^{(3)}{}_a \omega^{(3)}{}_b \right), \quad (7.36)$$

$$Q_{ab} = -r^2 \left(-\omega^{(0)}{}_a \omega^{(0)}{}_b + \omega^{(1)}{}_a \omega^{(1)}{}_b \right) + a^2 \cos^2 \theta \left(\omega^{(2)}{}_a \omega^{(2)}{}_b + \omega^{(3)}{}_a \omega^{(3)}{}_b \right). \quad (7.37)$$

We may use the expressions (7.34)–(7.37) to derive some useful identities for the Killing objects on Kerr spacetime. Notice from the form of (7.36) and (7.37) that

$$K_{ab} - Q_{ab} = (r^2 - a^2 \cos^2 \theta) g_{ab}, \quad (7.38)$$

where we have used (7.31). Now, contraction of (7.38) with $k^a k^b$ gives $K_{ab} k^a k^b - Q_{ab} k^a k^b = (r^2 - a^2 \cos^2 \theta) g_{ab} k^a k^b = 0$, which implies that

$$K = Q. \quad (7.39)$$

Using (7.31) and (7.36), we obtain two useful expressions for the Killing tensor:

$$K_{ab} = \Sigma \left(\omega^{(0)}{}_a \omega^{(0)}{}_b - \omega^{(1)}{}_a \omega^{(1)}{}_b \right) + r^2 g_{ab} \quad (7.40)$$

$$= \Sigma \left(\omega^{(2)}{}_a \omega^{(2)}{}_b + \omega^{(3)}{}_a \omega^{(3)}{}_b \right) - a^2 \cos^2 \theta g_{ab}. \quad (7.41)$$

Using (7.34) and (7.35) to read off the components of $f_{(\alpha)(\gamma)}$ and $h_{(\alpha)(\gamma)}$ (i.e., the projection of the conformal Killing–Yano tensor and the Killing–Yano tensor onto Carter's tetrad), one may check that

$$f_{(\alpha)(\gamma)} h_{(\beta)}{}^{(\gamma)} = f_{(\alpha)(\gamma)} \eta^{(\gamma)(\delta)} h_{(\delta)(\beta)} = ar \cos \theta \eta_{(\alpha)(\gamma)}, \quad (7.42)$$

where $\eta_{(\alpha)(\beta)}$ are the components of the Minkowski metric. Hence, the “product” of f_{ab} and h_{ab} is

$$f_{ac} h_b{}^c = h_{ac} f_b{}^c = ar \cos \theta g_{ab}. \quad (7.43)$$

7.3.4 Null geodesics: separability and conserved quantities

Let us consider an affinely parametrised null geodesic on Kerr spacetime with tangent vector field $k^a(s) = \dot{x}^a(s)$ in Boyer–Lindquist coordinates. The geodesic equation may be separated with the aid of the explicit and hidden symmetries of Kerr spacetime. The separation of geodesic motion was first demonstrated by Carter [47].

To demonstrate this separability, let us first consider the explicit symmetries which arise from the existence of the Killing vectors $\xi_{(t)}$ and $\xi_{(\phi)}$. In Boyer–Lindquist coordinates, we have $E = g_{ab}\xi_{(t)}^a k^b = -g_{tt}\dot{t} - g_{t\phi}\dot{\phi}$ and $L_z = g_{ab}\xi_{(\phi)}^a k^b = g_{t\phi}\dot{t} + g_{\phi\phi}\dot{\phi}$. These relationships can be inverted to give

$$\dot{t} = -g^{tt}E + g^{t\phi}L_z, \quad \dot{\phi} = -g^{t\phi}E + g^{\phi\phi}L_z. \quad (7.44)$$

Next, we may use (7.40) and (7.41) to separate out the r - and θ -motion. Contracting these identities with $k^a k^b$, we see that the terms involving g_{ab} will vanish due to the fact that k^a is null. We are therefore left with

$$K = \Sigma \left(\omega^{(0)}_a \omega^{(0)}_b - \omega^{(1)}_a \omega^{(1)}_b \right) k^a k^b \quad (7.45)$$

$$= \Sigma \left(\omega^{(2)}_a \omega^{(2)}_b + \omega^{(3)}_a \omega^{(3)}_b \right) k^a k^b. \quad (7.46)$$

Inserting the one-form components of Carter's orthonormal tetrad (7.23)–(7.26) into (7.45) and (7.46), and inserting the components of the inverse metric (7.14) and (7.15) into (7.44), we arrive at the geodesic equations

$$\dot{t} = \frac{1}{\Delta\Sigma} \left\{ E \left[(r^2 + a^2)^2 - \Delta a^2 \sin^2 \theta \right] - 2MarL_z \right\}, \quad (7.47)$$

$$\Sigma^2 \dot{r}^2 = \left[E (r^2 + a^2) - aL_z \right]^2 - \Delta K, \quad (7.48)$$

$$\Sigma^2 \dot{\theta}^2 = K - \left(aE \sin \theta - \frac{L_z}{\sin \theta} \right)^2, \quad (7.49)$$

$$\dot{\phi} = \frac{1}{\Delta\Sigma} \left[2MarE + \frac{(\Sigma - 2Mr)L_z}{\sin^2 \theta} \right]. \quad (7.50)$$

The constant K , which permits the separation of the radial and latitudinal motion, is the so-called *Carter constant*. This can be written using (7.48) or (7.49) as

$$K = \frac{1}{\Delta} \left\{ \left[E (r^2 + a^2) - aL_z \right]^2 - \Sigma^2 \dot{r}^2 \right\} = \Sigma^2 \dot{\theta}^2 + \left(aE \sin \theta - \frac{L_z}{\sin \theta} \right)^2. \quad (7.51)$$

Rather than work with the equations (7.47)–(7.50), we choose to employ the Hamiltonian formalism for null geodesics; see Section 2.2.3. The Hamiltonian takes the form $H = \frac{1}{2}g^{ab}p_a p_b$, where $p_a = g_{ab}\dot{x}^b = g_{ab}k^b$ are the conjugate momenta, and g^{ab} are the contravariant components of the metric tensor given by (7.14) and (7.15). The null geodesics are then the integral curves of Hamilton's equations,

$$\dot{x}^a = \frac{\partial H}{\partial p_a}, \quad \dot{p}_a = -\frac{\partial H}{\partial x^a}. \quad (7.52)$$

Clearly, the metric (7.13) is independent of the coordinates t and ϕ , so the momenta p_t and p_ϕ are constants of the motion. In fact, we have $E = -g_{ab}\xi_{(t)}^a k^b = -\xi_{(t)}^a p_a = -p_t$,

and $L_z = g_{ab}\xi_{(\phi)}^a k^b = \xi_{(\phi)}^a p_a = p_\phi$. The conserved quantities E and L_z correspond to the energy and azimuthal angular momentum of the photon, respectively. Since these conserved quantities are linear in the momenta, the symmetries generated by $\xi_{(t)}$ and $\xi_{(\phi)}$ are referred to as *explicit* symmetries [50].

As described in Section 7.3.2, there exist higher-rank Killing tensors on Kerr spacetime which give rise to further conserved quantities along rays. Here, we express these in terms of the Boyer–Lindquist coordinates and their conjugate momenta.

Firstly, we note that the metric itself is a (trivial) rank-two Killing tensor. The quantity $g^{ab}p_a p_b$ is therefore conserved along geodesics. This means that the Hamiltonian $H = \frac{1}{2}g^{ab}p_a p_b$ is a constant of the motion, with $H = 0$ in the case of null geodesics.

Furthermore, the Carter constant, which can be written in the form $K = K_{ab}k^a k^b$ is constant along geodesics. In terms of the Boyer–Lindquist coordinates and their conjugate momenta, the Carter constant can be expressed in the form

$$K = K^{ab}p_a p_b = p_\theta^2 + \left(aE \sin \theta - \frac{L_z}{\sin \theta} \right)^2. \quad (7.53)$$

Clearly, the Carter constant is quadratic in the momenta (by virtue of the fact that K_{ab} is a rank-two Killing tensor). The Carter constant therefore corresponds to a *hidden* symmetry of Kerr spacetime [50]. We also remark here that K is clearly non-negative. (One may arrive at an alternative expression for K by rearranging (7.48) and noting that $p_r = \frac{\Sigma}{\Delta}\dot{r}$.)

7.4 Construction of complex null tetrads

7.4.1 Complex null tetrads from symmetries of Kerr spacetime

We now use the Killing objects defined in Section 7.3.4 to construct a complex null tetrad $\{k^a, n^a, m^a, \bar{m}^a\}$, which is parallel-transported along a null geodesic with tangent vector field k^a . This is achieved by first constructing a non-parallel-transported tetrad $\{\tilde{k}^a, \tilde{n}^a, \tilde{m}^a, \tilde{\bar{m}}^a\}$, by projecting the tangent vector \tilde{k} , then performing a Lorentz transformation which keeps \tilde{k}^a fixed; see Section 2.5.2.

Define the complex bivector

$$\mathcal{F}_{ab} = f_{ab} - ih_{ab}, \quad (7.54)$$

where h_{ab} and f_{ab} are the closed conformal Killing–Yano tensor and the Killing–Yano tensor, respectively. Note that ${}^*\mathcal{F}_{ab} = -i\mathcal{F}_{ab}$, so \mathcal{F}_{ab} is self-dual. Using the properties of f_{ab} and h_{ab} , one may verify some useful identities for the complex bivector (7.54); see Appendix E. In particular, we have

$$\mathcal{F}_{ac}\mathcal{F}_b{}^c = (r - ia \cos \theta)^2 g_{ab}, \quad (7.55)$$

$$\mathcal{F}_{ac}\overline{\mathcal{F}}_b{}^c = K_{ab} + Q_{ab}. \quad (7.56)$$

We may now use \mathcal{F}_{ab} and its geometrical properties to construct the tetrad legs $\{\tilde{n}^a, \tilde{m}^a, \tilde{\bar{m}}^a\}$, by projecting the tangent vector \tilde{k}^a .

First, let

$$\tilde{m}^a = \frac{1}{\sqrt{2K}} \mathcal{F}^a_b \tilde{k}^b. \quad (7.57)$$

The leg \tilde{m}^a is obtained by complex conjugation of (7.57). Clearly \tilde{m}^a is orthogonal to \tilde{k}_a :

$$\tilde{k}^a \tilde{m}_a = \frac{1}{\sqrt{2K}} \tilde{k}^a \mathcal{F}_{ab} \tilde{k}^b = 0, \quad (7.58)$$

by virtue of the fact that \mathcal{F}_{ab} is antisymmetric ($\mathcal{F}_{ab} = -\mathcal{F}_{ba}$). Moreover, using the identity (7.55), we find

$$\tilde{m}^a \tilde{m}_a = \frac{1}{2K} \mathcal{F}^a_b \mathcal{F}_{ac} \tilde{k}^b \tilde{k}^c = \frac{1}{2K} (r - ia \cos \theta)^2 g_{bc} \tilde{k}^b \tilde{k}^c = 0, \quad (7.59)$$

i.e., \tilde{m}^a is null. Finally, using the identity (7.56), we see that

$$\tilde{m}^a \tilde{m}_a = \frac{1}{2K} \mathcal{F}^a_b \overline{\mathcal{F}}_{ac} \tilde{k}^b \tilde{k}^c = \frac{1}{2K} (K_{bc} + Q_{bc}) \tilde{k}^b \tilde{k}^c = 1, \quad (7.60)$$

where we recall that $K_{ab} \tilde{k}^a \tilde{k}^b = K = Q_{ab} \tilde{k}^a \tilde{k}^b$.

Applying the derivative operator $\tilde{D} = \tilde{k}^a \nabla_a$ to (7.57) gives (see Appendix E)

$$\tilde{D} \tilde{m}_a = \frac{iE}{\sqrt{2K}} \tilde{k}_a. \quad (7.61)$$

Clearly \tilde{m}^a is *not* parallel-propagated along a null geodesic with tangent vector \tilde{k}^a ; however, one can construct a parallel-transported vector m^a by performing a null rotation, as we shall demonstrate.

We define the final leg of our complex null tetrad to be

$$\tilde{n}^a = \frac{1}{2K} \mathcal{F}^{ab} \overline{\mathcal{F}}_{bc} \tilde{k}^c. \quad (7.62)$$

First, note that

$$\tilde{k}^a \tilde{n}_a = \frac{1}{2K} \mathcal{F}_a^b \overline{\mathcal{F}}_{bc} \tilde{k}^a \tilde{k}^c = -\frac{1}{2K} (K_{ac} + Q_{ac}) \tilde{k}^a \tilde{k}^c = -1, \quad (7.63)$$

where we have used (7.56) and the fact that $K_{ab} \tilde{k}^a \tilde{k}^b = K = Q_{ab} \tilde{k}^a \tilde{k}^b$. Next, taking the inner product with \tilde{m}^a , we see that

$$\tilde{n}^a \tilde{m}_a = \frac{1}{(2K)^{3/2}} \mathcal{F}^{ab} \overline{\mathcal{F}}_{bc} \mathcal{F}_{ad} \tilde{k}^c \tilde{k}^d = \frac{1}{(2K)^{3/2}} (r - ia \cos \theta)^2 \overline{\mathcal{F}}_{cd} \tilde{k}^c \tilde{k}^d = 0, \quad (7.64)$$

where we have used (7.55) and the fact that $\overline{\mathcal{F}}_{ab}$ is antisymmetric. (It follows from complex conjugation that $\tilde{n}^a \tilde{m}_a = 0$, since \tilde{n}^a is real.) Finally,

$$\tilde{n}^a \tilde{n}_a = \frac{1}{(2K)^2} \mathcal{F}^{ab} \overline{\mathcal{F}}_{bc} \tilde{k}^c \mathcal{F}_a^d \overline{\mathcal{F}}_{de} \tilde{k}^e = \left(\frac{\Sigma}{2K} \right)^2 g_{ce} \tilde{k}^c \tilde{k}^e = 0, \quad (7.65)$$

where we have used the identity (7.55) and its complex conjugate, as well as the fact that \tilde{k}^a is null. Hence, \tilde{n}^a is null. We note that \tilde{n}^a is *not* parallel-propagated along null geodesics.

We have used the complex self-dual bivector (7.54) to construct a complex null tetrad $\{\tilde{k}^a, \tilde{n}^a, \tilde{m}^a, \overline{\tilde{m}^a}\}$, with all inner products zero except $\tilde{k}^a \tilde{n}_a = -1$ and $\tilde{m}^a \overline{\tilde{m}^a} = 1$. Only

the tangent vector \tilde{k}^a is parallel-transported. However, we may construct a parallel-transported tetrad by performing a Lorentz transformation of the form

$$k^a = \tilde{k}^a, \quad m^a = \tilde{m}^a + B\tilde{k}^a, \quad n^a = \tilde{n}^a + B\tilde{\bar{m}}^a + \bar{B}\tilde{m}^a + B\bar{B}\tilde{k}^a. \quad (7.66)$$

We note that the transformation (7.66) preserves the inner product relationships between the tetrad legs, so that $k^a n_a = -1$ and $m^a \bar{m}_a = 1$, with all others zero.

Clearly k^a is parallel-transported. The requirement that m^a be parallel-transported fixes the null rotation parameter B in (7.66) as

$$B = -\frac{iEs}{\sqrt{2K}} + B_0, \quad (7.67)$$

where $B_0 \in \mathbb{C}$ is an arbitrary integration constant. The transport equation (7.61) then reads

$$D\tilde{m}_a = -\dot{B}k_a. \quad (7.68)$$

It follows from the fact that each of the legs $\{k^a, m^a, \bar{m}^a\}$ are parallel-transported that n^a is also parallel-propagated along null geodesics.

In summary, taking the complex null tetrad $\{\tilde{k}^a, \tilde{n}^a, \tilde{m}^a, \tilde{\bar{m}}^a\}$, and performing the Lorentz transformation (7.66) with null rotation parameter (7.67) yields a complex null tetrad $\{k^a, n^a, m^a, \bar{m}^a\}$, which is parallel-transported along null geodesics.

Intriguingly, we can absorb the Lorentz transformation (7.66) into the self-dual bivector (7.54) by defining a new “projection operator”

$$\mathcal{G}_{ab} = \mathcal{F}_{ab} + \sqrt{2K}B g_{ab}, \quad (7.69)$$

where B is defined in (7.67). (Note that \mathcal{G}_{ab} is not a bivector.) Then the definitions (7.57) and (7.62) with \mathcal{F}_{ab} replaced by \mathcal{G}_{ab} (and “tilded” quantities replaced by “untilded” quantities) yield the parallel-transported legs m^a and n^a . The procedure outlined above applies to all geometries which admit a principal tensor, i.e., the (off-shell) Kerr–NUT–anti-de Sitter family of spacetimes [50].

We caution here that the construction of these tetrads is valid only for rays with $K > 0$. Constructing a parallel-propagated complex null tetrad along rays with $K = 0$ requires a special treatment.

On Kerr spacetime, the “tilded” complex null tetrad derived here can be related to the well-known Kinnersley tetrad by application of a sequence of Lorentz transformations. This is done explicitly in Appendix E.

7.4.2 Marck’s tetrads

Quasi-orthonormal tetrad

Marck [52] employs Carter’s canonical tetrad (Section 7.3.3) and the symmetries of Kerr spacetime to construct a tetrad which is parallel-propagated along null geodesics in the Kerr(–Newman) spacetime.

This is achieved by first constructing a quasi-orthonormal tetrad $\{\tilde{\lambda}_0^a, \tilde{\lambda}_1^a, \tilde{\lambda}_2^a, \tilde{\lambda}_3^a\}$, where $\tilde{\lambda}_0^a$ and $\tilde{\lambda}_3^a$ are null, and $\tilde{\lambda}_1^a$ and $\tilde{\lambda}_2^a$ are unit spacelike. Marck identifies $\tilde{\lambda}_0^a$ with

the tangent vector k^a , which is clearly parallel-transported. The components of this tetrad leg with respect to Carter's symmetric tetrad $\tilde{\lambda}_0^{(\alpha)} = \tilde{\lambda}_0^a \omega^{(\alpha)}_a$ are given by

$$\tilde{\lambda}_0^{(0)} = \frac{1}{\sqrt{\Delta\Sigma}} [E(r^2 + a^2) - aL_z], \quad (7.70)$$

$$\tilde{\lambda}_0^{(1)} = \sqrt{\frac{\Sigma}{\Delta}} \dot{r}, \quad (7.71)$$

$$\tilde{\lambda}_0^{(2)} = \sqrt{\Sigma} \dot{\theta}, \quad (7.72)$$

$$\tilde{\lambda}_0^{(3)} = \frac{1}{\sqrt{\Sigma}} \left(aE \sin \theta - \frac{L_z}{\sin \theta} \right). \quad (7.73)$$

A second parallel-transported vector which is orthogonal to $\tilde{\lambda}_0^{(\alpha)}$ arises naturally due to the existence of the non-trivial Killing–Yano tensor $f_{(\alpha)(\beta)}$. The properties of this tensor (discussed in Section 7.3.2) ensure that the vector with components $v^{(\alpha)} = f^{(\alpha)}_{(\beta)} \tilde{\lambda}_0^{(\beta)}$ is (i) parallel-propagated along the geodesic with tangent vector $\tilde{\lambda}_0^{(\alpha)}$, and (ii) orthogonal to $\tilde{\lambda}_0^{(\alpha)}$. The leg $\tilde{\lambda}_2^{(\alpha)}$ is chosen to be the unit vector parallel to $v^{(\alpha)}$; its components with respect to the symmetric tetrad are

$$\tilde{\lambda}_2^{(0)} = \sqrt{\frac{\Sigma}{K\Delta}} a \cos \theta \dot{r} = \frac{a \cos \theta}{\sqrt{K}} \tilde{\lambda}_0^{(1)}, \quad (7.74)$$

$$\tilde{\lambda}_2^{(1)} = \frac{a \cos \theta}{\sqrt{K\Delta\Sigma}} [E(r^2 + a^2) - aL_z] = \frac{a \cos \theta}{\sqrt{K}} \tilde{\lambda}_0^{(0)}, \quad (7.75)$$

$$\tilde{\lambda}_2^{(2)} = -\frac{r}{\sqrt{K\Sigma}} \left(aE \sin \theta - \frac{L_z}{\sin \theta} \right) = -\frac{r}{\sqrt{K}} \tilde{\lambda}_0^{(3)}, \quad (7.76)$$

$$\tilde{\lambda}_2^{(3)} = \sqrt{\frac{\Sigma}{K}} r \dot{\theta} = \frac{r}{\sqrt{K}} \tilde{\lambda}_0^{(2)}. \quad (7.77)$$

There is now a straightforward choice for the vectors $\tilde{\lambda}_1^a$ and $\tilde{\lambda}_3^a$, such that the tetrad legs satisfy

$$\tilde{\lambda}_0^{(\alpha)} \tilde{\lambda}_{3(\alpha)} = \tilde{\lambda}_1^{(\alpha)} \tilde{\lambda}_{1(\alpha)} = \tilde{\lambda}_2^{(\alpha)} \tilde{\lambda}_{2(\alpha)} = 1, \quad (7.78)$$

with all other inner products zero. This choice is given by

$$\tilde{\lambda}_1^{(0)} = \sqrt{\frac{\Sigma}{K\Delta}} r \dot{r} = \frac{r}{\sqrt{K}} \tilde{\lambda}_0^{(1)}, \quad (7.79)$$

$$\tilde{\lambda}_1^{(1)} = \frac{r}{\sqrt{K\Delta\Sigma}} [E(r^2 + a^2) - aL_z] = \frac{r}{\sqrt{K}} \tilde{\lambda}_0^{(0)}, \quad (7.80)$$

$$\tilde{\lambda}_1^{(2)} = \frac{a \cos \theta}{\sqrt{K\Sigma}} \left(aE \sin \theta - \frac{L_z}{\sin \theta} \right) = \frac{a \cos \theta}{\sqrt{K}} \tilde{\lambda}_0^{(3)}, \quad (7.81)$$

$$\tilde{\lambda}_1^{(3)} = \sqrt{\frac{\Sigma}{K}} a \cos \theta \dot{\theta} = -\frac{a \cos \theta}{\sqrt{K}} \tilde{\lambda}_0^{(2)}, \quad (7.82)$$

and

$$\tilde{\lambda}_3^{(0)} = -\frac{1}{2K} \sqrt{\frac{\Sigma}{\Delta}} [E(r^2 + a^2) - aL_z] = -\frac{\Sigma}{2K} \tilde{\lambda}_0^{(0)}, \quad (7.83)$$

$$\tilde{\lambda}_3^{(1)} = -\frac{1}{2K} \sqrt{\frac{\Sigma^3}{\Delta}} \dot{r} = -\frac{\Sigma}{2K} \tilde{\lambda}_0^{(1)}, \quad (7.84)$$

$$\tilde{\lambda}_3^{(2)} = \frac{1}{2K} \sqrt{\Sigma^3} \dot{\theta} = \frac{\Sigma}{2K} \tilde{\lambda}_0^{(2)}, \quad (7.85)$$

$$\tilde{\lambda}_3^{(3)} = \frac{1}{2K} \sqrt{\Sigma} \left(aE \sin \theta - \frac{L_z}{\sin \theta} \right) = \frac{\Sigma}{2K} \tilde{\lambda}_0^{(3)}. \quad (7.86)$$

Unlike $\tilde{\lambda}_0^a$ and $\tilde{\lambda}_2^a$, the legs $\tilde{\lambda}_1^a$ and $\tilde{\lambda}_3^a$ are *not* parallel-propagated along null geodesics with tangent vector $\tilde{\lambda}_0^a$.

Parallel-transported tetrad

Marck [52] demonstrates that one may use the tetrad $\{\tilde{\lambda}_0^a, \tilde{\lambda}_1^a, \tilde{\lambda}_2^a, \tilde{\lambda}_3^a\}$ to construct a parallel-transported tetrad $\{\lambda_0^a, \lambda_1^a, \lambda_2^a, \lambda_3^a\}$. First, one may set $\lambda_0^a = \tilde{\lambda}_0^a$ and $\lambda_2^a = \tilde{\lambda}_2^a$, as these legs are already parallel-propagated. The remaining legs, λ_1^a and λ_3^a , may be constructed by means of a one-parameter Lorentz transformation of the form

$$\lambda_1^a = \tilde{\lambda}_1^a - \Psi \tilde{\lambda}_0^a, \quad \lambda_3^a = \tilde{\lambda}_3^a + \Psi \tilde{\lambda}_1^a - \frac{1}{2} \Psi^2 \tilde{\lambda}_0^a, \quad (7.87)$$

where Ψ is a real-valued function. The requirement that the legs be parallel-transported yields a first-order ordinary differential equation for Ψ which reads $\dot{\Psi} = \left(k^b \nabla_b \tilde{\lambda}_1^a \right) \tilde{\lambda}_{3a}$. Integration of this equation permits us to express the parameter Ψ as [52]

$$\Psi = \frac{1}{\sqrt{K}} \left(\int \frac{E(r^2 + a^2) - aL_z}{[R(r)]^{1/2}} dr - \int a \sin \theta \frac{aE \sin \theta - L_z \operatorname{cosec} \theta}{[\Theta(\theta)]^{1/2}} d\theta \right). \quad (7.88)$$

Here, $R(r) = [E(r^2 + a^2) - aL_z]^2 - K\Delta$ and $\Theta(\theta) = K - (aE \sin \theta - L_z \operatorname{cosec} \theta)^2$, where the signs of $R^{1/2} = \pm\sqrt{R}$ and $\Theta^{1/2} = \pm\sqrt{\Theta}$ are chosen independently and in accordance with the turning points in the r - and θ -motion.

Inserting the equations of motion (7.48) and (7.49) into (7.88), one can simplify the expression for the tetrad parameter Ψ :

$$\Psi = \frac{1}{\sqrt{K}} \int \left(\frac{E(r^2 + a^2) - aL_z}{[R(r)]^{1/2}} \dot{r} - a \sin \theta \frac{aE \sin \theta - L_z \operatorname{cosec} \theta}{[\Theta(\theta)]^{1/2}} \dot{\theta} \right) ds \quad (7.89)$$

$$= \frac{1}{\sqrt{K}} \int \frac{E(r^2 + a^2 \cos^2 \theta)}{\Sigma} ds \quad (7.90)$$

$$= \frac{E}{\sqrt{K}} \int ds. \quad (7.91)$$

Hence, the Lorentz transformation parameter Ψ takes the form

$$\Psi = \frac{Es}{\sqrt{K}} + \psi, \quad (7.92)$$

where s denotes the affine parameter and ψ is an arbitrary constant of integration. One can see from (7.53) that the Carter constant is non-negative. This construction of the parallel-propagated tetrad is therefore valid for all null rays, apart from those which satisfy $K = 0$. In this degenerate case, the vector $v^{(\alpha)} = f^{(\alpha)}_{(\beta)} \tilde{\lambda}_0^{(\beta)}$ is parallel to $\lambda_0^{(\alpha)}$; a special treatment would therefore be required to construct a parallel-propagated tetrad along null rays with vanishing Carter constant.

Complex null tetrads from Marck's tetrads

One may construct a complex null tetrad from Marck's quasi-orthonormal tetrad via

$$\tilde{k}^a = \tilde{\lambda}_0^a, \quad \tilde{n}^a = -\tilde{\lambda}_3^a, \quad \tilde{m}^a = \frac{1}{\sqrt{2}} \left(\tilde{\lambda}_2^a + i\tilde{\lambda}_1^a \right), \quad \tilde{\bar{m}}^a = \frac{1}{\sqrt{2}} \left(\tilde{\lambda}_2^a - i\tilde{\lambda}_1^a \right). \quad (7.93)$$

This tetrad is a Newman–Penrose tetrad: the legs satisfy $\tilde{k}^a \tilde{n}_a = -1$ and $\tilde{m}^a \tilde{\bar{m}}_a = 1$, with all other inner products zero. We note that only \tilde{k}^a is parallel-transported.

One may construct a parallel-transported complex null tetrad by replacing the “tilded” legs with parallel-transported legs ($\tilde{\lambda}_i^a \mapsto \lambda_i^a$) in the definitions (7.93). Alternatively, one may perform a null rotation of the tetrad (7.93) of the form (7.66). The requirement that the tetrad $\{k^a, n^a, m^a, \bar{m}^a\}$ be parallel-transported yields $\dot{B} = -\tilde{n}^a D \tilde{m}_a$. Integrating this ordinary differential equation fixes the null rotation parameter:

$$B = -\frac{i\Psi}{\sqrt{2}}, \quad (7.94)$$

where Ψ is Marck’s tetrad parameter, given in (7.92).

It should be noted that the complex null tetrads built from Marck’s tetrads are equivalent to those which were constructed using the complex self-dual bivector \mathcal{F}_{ab} in Section 7.4.1. This is best seen by projecting onto Carter’s symmetric tetrad, as shown in Appendix E. Moreover, direct substitution of Marck’s tetrad parameter (7.92) into (7.94) shows that the null rotation parameters (7.67) and (7.94) are equivalent.

7.5 Newman–Penrose formalism and transport equations

7.5.1 Transport equations for Newman–Penrose quantities

The Newman–Penrose scalars are defined in terms of projections of first derivatives of the complex null tetrad legs [93]; see Section 2.5 for a review. We define directional derivatives along the tetrad legs to be

$$D = k^a \nabla_a, \quad \Delta = n^a \nabla_a, \quad \delta = m^a \nabla_a, \quad \bar{\delta} = \bar{m}^a \nabla_a. \quad (7.95)$$

For our parallel-propagated complex null tetrad (Section 7.4), three of the Newman–Penrose scalars are trivially zero: $\kappa = \pi = \epsilon = 0$. The eight non-zero Newman–Penrose scalars which we use are

$$\rho = -m^a k_{a;b} \bar{m}^b, \quad \sigma = -m^a k_{a;b} m^b, \quad (7.96)$$

$$\chi = -\bar{m}^a m_{a;b} m^b, \quad \tau = -m^a k_{a;b} n^b, \quad (7.97)$$

$$\mu = \bar{m}^a n_{a;b} m^b, \quad \lambda = \bar{m}^a n_{a;b} \bar{m}^b, \quad (7.98)$$

$$\nu = \bar{m}^a n_{a;b} n^b, \quad \gamma = \frac{1}{2} \left(\bar{m}^a m_{a;b} n^b - n^a k_{a;b} n^b \right). \quad (7.99)$$

One can obtain identities for the Newman–Penrose scalars by using the fact that $g^{ab} = -k^a n^b - n^a k^b + m^a \bar{m}^b + \bar{m}^a m^b$, and using the fact that $k_a = \Phi_{,a}$ is a gradient. For example, it is straightforward to see that the twist-free property of the complex null tetrad (i.e., $k_{[a;b]} = 0$) implies that ρ is a real quantity. In fact, $\rho = -\frac{1}{2}\vartheta$, where $\vartheta = k^a_{;a}$ is the expansion scalar [172]. For our tetrad, $\tau = \beta + \bar{\alpha}$ and $\chi = \beta - \bar{\alpha}$, where $\alpha = \frac{1}{2}(k^a n_{a;b} \bar{m}^b - m^a \bar{m}_{a;b} m^b)$ and $\beta = \frac{1}{2}(\bar{m}^a m_{a;b} m^b - n^a k_{a;b} n^b)$. It is therefore possible to eliminate both α and β with the introduction of χ [81].

Along null geodesics, the Newman–Penrose scalars satisfy a system of transport equations [83]. In a Ricci-flat spacetime, the transport equations we require are

$$D\rho = \rho^2 + \sigma\bar{\sigma}, \quad D\sigma = 2\rho\sigma + \Psi_0, \quad (7.100)$$

$$D\chi = \rho\chi - \sigma\bar{\chi} + \Psi_1, \quad D\tau = \rho\tau + \sigma\bar{\tau} + \Psi_1, \quad (7.101)$$

$$D\mu = \rho\mu + \sigma\lambda + \Psi_2, \quad D\lambda = \rho\lambda + \bar{\sigma}\mu, \quad (7.102)$$

$$D\nu = \bar{\tau}\mu + \tau\lambda + \Psi_3, \quad D\gamma = \tau\bar{\tau} + \frac{1}{2}(\bar{\tau}\chi - \tau\bar{\chi}) + \Psi_2. \quad (7.103)$$

Here, Ψ_i are the complex Weyl scalars (see Section 2.5.4). We calculate these for the “tilded” (non-parallel-transported) tetrad and the parallel-transported tetrad in Section 7.5.2. In this work, we are interested the closed subsystem of equations (7.100)–(7.102), satisfied by the Newman–Penrose scalars $\{\rho, \sigma, \chi, \tau, \mu, \lambda\}$.

In addition to the Newman–Penrose scalars, we wish to calculate higher-order quantities, such as $\delta\mathcal{A}^2$ and $\delta\bar{\chi}$, which feature in the transport equations for sub-leading-order geometrical optics quantities (e.g. \mathbf{u}_1); see Section 7.2.2 and [81]. In general, one may obtain transport equations by using the commutator identity $D\delta = \delta D - \tau D + \rho\delta + \sigma\bar{\delta}$ and its complex conjugate [83]. The system of transport equations we require is

$$D(\delta\mathcal{A}^2) = 3\rho\delta\mathcal{A}^2 + \sigma\bar{\delta}\mathcal{A}^2 + 2\mathcal{A}^2(\delta\rho - \rho\tau), \quad (7.104)$$

$$D(\bar{\delta}\mathcal{A}^2) = 3\rho\bar{\delta}\mathcal{A}^2 + \sigma\delta\mathcal{A}^2 + 2\mathcal{A}^2(\bar{\delta}\rho - \rho\bar{\tau}), \quad (7.105)$$

$$D(\delta\rho) = 3\rho\delta\rho + \bar{\sigma}\delta\sigma + \sigma(\bar{\delta}\sigma)^* - (\rho^2 + \sigma\bar{\sigma})\tau + \sigma\bar{\delta}\rho, \quad (7.106)$$

$$D(\bar{\delta}\rho) = 3\rho\bar{\delta}\rho + \bar{\sigma}\bar{\delta}\sigma + \sigma(\delta\sigma)^* - (\rho^2 + \sigma\bar{\sigma})\bar{\tau} + \bar{\sigma}\delta\rho, \quad (7.107)$$

$$D(\delta\sigma) = 3\rho\delta\sigma + 2\sigma\bar{\delta}\rho - 2\rho\sigma\tau + \sigma\bar{\delta}\sigma - \tau\Psi_0 + \delta\Psi_0, \quad (7.108)$$

$$D(\bar{\delta}\sigma) = 3\rho\bar{\delta}\sigma + 2\sigma\bar{\delta}\rho - 2\rho\sigma\bar{\tau} + \bar{\sigma}\delta\sigma - \bar{\tau}\Psi_0 + \bar{\delta}\Psi_0, \quad (7.109)$$

$$D(\delta\chi) = 2\rho\delta\chi + \chi\delta\rho + \sigma\bar{\delta}\chi - \sigma(\bar{\delta}\chi)^* - \bar{\chi}\delta\sigma - \tau(\rho\chi - \sigma\bar{\chi} + \Psi_1) + \delta\Psi_1, \quad (7.110)$$

$$D(\bar{\delta}\chi) = 2\rho\bar{\delta}\chi + \chi\bar{\delta}\rho + \bar{\sigma}\delta\chi - \sigma(\delta\chi)^* - \bar{\chi}\bar{\delta}\sigma - \bar{\tau}(\rho\chi - \sigma\bar{\chi} + \Psi_1) + \bar{\delta}\Psi_1, \quad (7.111)$$

where for extra clarity we use an asterisk (as well as an overbar) to denote complex conjugation. This system of transport equations features directional derivatives of the first two Weyl scalars along the legs m and \bar{m} : $\{\delta\Psi_0, \bar{\delta}\Psi_0, \delta\Psi_1, \bar{\delta}\Psi_1\}$. These quantities are calculated in Section 7.5.3.

7.5.2 Weyl scalars

The Weyl tensor C_{abcd} , which in this case is identical to the Riemann tensor due to the fact that Kerr spacetime is Ricci-flat, can be expressed in terms of Carter’s symmetric

tetrad as [51, 315]

$$\Omega_{(1)}^{(0)} = 2I_1 \omega^{(0)} \wedge \omega^{(1)} + 2I_2 \omega^{(2)} \wedge \omega^{(3)}, \quad (7.112)$$

$$\Omega_{(2)}^{(0)} = -I_1 \omega^{(0)} \wedge \omega^{(2)} + I_2 \omega^{(1)} \wedge \omega^{(3)}, \quad (7.113)$$

$$\Omega_{(3)}^{(0)} = -I_1 \omega^{(0)} \wedge \omega^{(3)} - I_2 \omega^{(1)} \wedge \omega^{(2)}, \quad (7.114)$$

$$\Omega_{(2)}^{(1)} = -I_1 \omega^{(1)} \wedge \omega^{(2)} + I_2 \omega^{(0)} \wedge \omega^{(3)}, \quad (7.115)$$

$$\Omega_{(1)}^{(3)} = I_1 \omega^{(1)} \wedge \omega^{(3)} + I_2 \omega^{(0)} \wedge \omega^{(2)}, \quad (7.116)$$

$$\Omega_{(2)}^{(3)} = -2I_1 \omega^{(2)} \wedge \omega^{(3)} + 2I_2 \omega^{(0)} \wedge \omega^{(1)}, \quad (7.117)$$

where $\Omega_{(\beta)}^{(\alpha)} = \frac{1}{2} C_{(\beta)(\gamma)(\delta)}^{(\alpha)} \omega^{(\gamma)} \wedge \omega^{(\delta)}$ is the *curvature two-form*, and the functions I_1 and I_2 are defined to be

$$I_1 = \frac{Mr}{\Sigma^3} (r^2 - 3a^2 \cos^2 \theta), \quad I_2 = \frac{Ma \cos \theta}{\Sigma^3} (3r^2 - a^2 \cos^2 \theta). \quad (7.118)$$

These quantities satisfy the identity $I_1 + iI_2 = \frac{M}{(r - ia \cos \theta)^3}$.

Using the components of the Weyl tensor in terms of Carter's symmetric tetrad, we calculate the complex Weyl scalars in the “tilded” complex null tetrad. We find

$$\tilde{\Psi}_0 = C_{\tilde{k}\tilde{m}\tilde{k}\tilde{m}} = \frac{3KM}{(r + ia \cos \theta)^5}, \quad (7.119)$$

$$\tilde{\Psi}_1 = C_{\tilde{k}\tilde{n}\tilde{k}\tilde{m}} = 0, \quad (7.120)$$

$$\tilde{\Psi}_2 = C_{\tilde{k}\tilde{m}\tilde{m}\tilde{n}} = \frac{M}{2(r + ia \cos \theta)^3}, \quad (7.121)$$

$$\tilde{\Psi}_3 = C_{\tilde{k}\tilde{n}\tilde{m}\tilde{n}} = 0, \quad (7.122)$$

$$\tilde{\Psi}_4 = C_{\tilde{m}\tilde{n}\tilde{m}\tilde{n}} = \frac{3M}{4K(r + ia \cos \theta)}, \quad (7.123)$$

where we employ the notation $C_{\tilde{k}\tilde{m}\tilde{k}\tilde{m}} = C_{abcd} \tilde{k}^a \tilde{m}^b \tilde{k}^c \tilde{m}^d$ for brevity.

Recall that the “tilded” complex null tetrad $\{\tilde{k}, \tilde{n}, \tilde{m}, \tilde{\bar{m}}\}$ and the parallel-transported complex null tetrad $\{k, n, m, \bar{m}\}$ are related by the Lorentz transformation (7.66) with (7.67). The Weyl scalars in the parallel-propagated frame can be expressed in terms of (7.119)–(7.123) as [83]

$$\Psi_0 = C_{kmkm} = \tilde{\Psi}_0, \quad (7.124)$$

$$\Psi_1 = C_{knkm} = \tilde{\Psi}_1 + \bar{B} \tilde{\Psi}_0, \quad (7.125)$$

$$\Psi_2 = C_{km\bar{m}n} = \tilde{\Psi}_2 + 2\bar{B} \tilde{\Psi}_1 + \bar{B}^2 \tilde{\Psi}_0, \quad (7.126)$$

$$\Psi_3 = C_{kn\bar{m}n} = \tilde{\Psi}_3 + 3\bar{B} \tilde{\Psi}_2 + 3\bar{B}^2 \tilde{\Psi}_1 + \bar{B}^3 \tilde{\Psi}_0, \quad (7.127)$$

$$\Psi_4 = C_{\bar{m}n\bar{m}n} = \tilde{\Psi}_4 + 4\bar{B} \tilde{\Psi}_3 + 6\bar{B}^2 \tilde{\Psi}_2 + 4\bar{B}^3 \tilde{\Psi}_1 + \bar{B}^4 \tilde{\Psi}_0, \quad (7.128)$$

with B given by (7.67).

7.5.3 Directional derivatives of Weyl scalars

The transport equations in the higher-order geometric optics formalism feature the quantities $\{\delta\Psi_0, \bar{\delta}\Psi_0, \delta\Psi_1, \bar{\delta}\Psi_1\}$, i.e., directional derivatives along m and \bar{m} of the Weyl scalars

Ψ_0 and Ψ_1 . The calculation of these quantities is more involved than the calculation of the Weyl scalars.

First, note that the Weyl scalars of interest are $\Psi_0 = C_{kmkm} = C_{abcd}k^am^bk^cm^d$ and $\Psi_1 = C_{km\bar{m}m} = C_{abcd}k^am^b\bar{m}^cm^d$. Applying the differential operators δ and $\bar{\delta}$ to these quantities and using the Leibniz rule, we see that we will require directional derivatives (along m and \bar{m}) of the tetrad legs k , m and \bar{m} . These are given by

$$\delta k^a = \tau k^a - \rho m^a - \sigma \bar{m}^a, \quad \bar{\delta} k^a = \bar{\tau} k^a - \bar{\sigma} m^a - \rho \bar{m}^a, \quad (7.129)$$

$$\delta m^a = \bar{\lambda} k^a - \sigma n^a + \chi m^a, \quad \bar{\delta} m^a = \lambda k^a - \bar{\sigma} n^a + \bar{\chi} \bar{m}^a. \quad (7.130)$$

The directional derivatives of the Weyl scalars can therefore be obtained by applying the derivative operators δ and $\bar{\delta}$ to the expressions $\Psi_0 = C_{abcd}k^am^bk^cm^d$ and $\Psi_1 = C_{abcd}k^am^b\bar{m}^cm^d$, and using the identities (7.129)–(7.130). This yields

$$\delta \Psi_0 = 2(\tau + \chi)\Psi_0 - 4\sigma\Psi_1 + (\delta C)_{kmkm}, \quad (7.131)$$

$$\bar{\delta} \Psi_0 = 2(\bar{\tau} - \bar{\chi})\Psi_0 - 4\sigma\Psi_1 + (\bar{\delta} C)_{kmkm}, \quad (7.132)$$

$$\delta \Psi_1 = \mu\Psi_0 + (\tau + \chi)\Psi_1 - 3\sigma\Psi_2 + (\delta C)_{km\bar{m}m}, \quad (7.133)$$

$$\bar{\delta} \Psi_1 = \lambda\Psi_0 + (\bar{\tau} - \bar{\chi})\Psi_1 - 3\rho\Psi_2 + (\bar{\delta} C)_{km\bar{m}m}, \quad (7.134)$$

where $(\delta C)_{kmkm} = (m^e \nabla_e C_{abcd})k^am^bk^cm^d$, for example. The quantities $(\delta C)_{kmkm}$, $(\bar{\delta} C)_{kmkm}$, $(\delta C)_{km\bar{m}m}$ and $(\bar{\delta} C)_{km\bar{m}m}$ can be calculated explicitly by transforming between the Kinnersley tetrad and the parallel-transported complex null tetrad; this calculation is presented in Appendix E.

7.5.4 Constraints from Newman–Penrose field equations

When projected onto a complex null tetrad, the Ricci identities give rise to a system of Newman–Penrose field equations, which relate the Newman–Penrose scalars, the Weyl scalars and the Ricci scalars; the full system of equations can be found in (e.g.) [83]. The Newman–Penrose field equations relevant for our purposes are

$$Z = \delta\rho - \bar{\delta}\sigma - \rho\tau + \sigma(\bar{\tau} - 2\bar{\chi}) + \Psi_1 = 0, \quad (7.135)$$

$$W = \bar{\delta}\chi + (\bar{\delta}\chi)^* + \rho(\mu + \bar{\mu}) - \sigma\lambda - \bar{\sigma}\bar{\lambda} + 2\chi\bar{\chi} - \Psi_2 - \bar{\Psi}_2 = 0. \quad (7.136)$$

We note that W is real. The identities (7.135) and (7.136) can be viewed as algebraic constraints which must be satisfied along null rays by the Newman–Penrose scalars, higher-order Newman–Penrose quantities, and Weyl curvature scalars.

With some work, it is possible to use the transport equations of Section 7.5.1 to demonstrate that the quantities Z and W satisfy a coupled system of transport equations along rays, namely

$$DZ = 3\rho Z - \bar{\sigma}Z, \quad DW = 2\rho W + \bar{\chi}Z + \chi\bar{Z}. \quad (7.137)$$

These equations indicate that, if $Z \neq 0$ or $W \neq 0$ initially (perhaps due to numerical error), then the quantities Z and W will grow, which could cause propagation of errors when numerically integrating transport equations along rays. It is therefore important to keep track of the constraints (i.e., Ricci identities) along rays.

7.6 Far-field asymptotics

We are interested in the special case of an initially parallel null geodesic congruence which starts at $r = \infty$. In practice – i.e., when numerically integrating our transport equations – we must start at some finite radius $r = r_0$. We therefore wish to determine the far-field behaviour of the Newman–Penrose scalars and higher-order quantities as power series in r^{-1} .

We first note that, for $r \rightarrow \infty$, the geodesic equation (7.48) is $\dot{r}^2 \sim E^2$. This has solution $r \sim \pm Es + r_0$, where $r_0 = r(0)$. Here, the upper (+) sign corresponds to *outgoing* rays (with $\dot{r} > 0$), and the lower (–) sign corresponds to *ingoing* rays (with $\dot{r} < 0$). We consider the latter case, as we are interested in the far-field behaviour of the system *before* scattering has occurred.

Recall that the Weyl scalars in the parallel-transported basis are related to those of the “tilded” tetrad by a parabolic Lorentz transformation, which features the complex parameter (7.67), which can be rewritten in the form

$$B = -\frac{i(Es + b_0)}{\sqrt{2K}}, \quad (7.138)$$

where $b_0 \in \mathbb{C}$ is an arbitrary integration constant. For rays which are ingoing from $r = r_0 \gg 1$, we have $Es \sim r_0 - r$. Making the replacement $Es \mapsto r - r_0$, and choosing $b_0 = -r_0$ to eliminate any dependence on r_0 in (7.138) yields

$$B = \frac{ir}{\sqrt{2K}}. \quad (7.139)$$

7.6.1 Far-field behaviour of the Weyl scalars

Recall the Weyl scalars in the parallel-propagated frame (7.124)–(7.128). Inserting the far-field null rotation parameter (7.139) into these expressions, and expanding to leading order in r^{-1} , we find

$$\Psi_0 = 3MKr^{-5} + O(r^{-6}), \quad (7.140)$$

$$\Psi_1 = -3iM\sqrt{K}r^{-4} + O(r^{-5}), \quad (7.141)$$

$$\Psi_2 = -Mr^{-3} + O(r^{-4}), \quad (7.142)$$

$$\Psi_3 = \frac{3Ma \cos \theta}{\sqrt{2K}}r^{-3} + O(r^{-4}), \quad (7.143)$$

$$\Psi_4 = -\frac{3Ma^2 \cos^2 \theta}{K}r^{-3} + O(r^{-4}). \quad (7.144)$$

We see that the first three Weyl scalars (7.140)–(7.142) exhibit a peeling behaviour at infinity: $\Psi_i = O(r^{i-5})$ for $i \in \{0, 1, 2\}$. Moreover, the leading-order terms for these quantities are independent of the black hole spin a . The leading-order solutions for the remaining Weyl scalars (7.143) and (7.144) depend on the spin of the black hole a , and the polar angle θ . The Weyl scalars Ψ_3 and Ψ_4 therefore exhibit a first-order post-Newtonian effect.

7.6.2 Far-field behaviour of the Newman–Penrose scalars

The far-field behaviour of the Newman–Penrose scalars may be determined using the transport equations (7.100)–(7.102), and the far-field expressions for the Weyl scalars (7.140)–(7.144). We note that, in the far-field limit, the derivative operator which appears on the left-hand side of the transport equations may be recast in terms of the radial coordinate using the chain rule: $D = \frac{d}{ds} = \dot{r} \frac{d}{dr} \sim -E \frac{d}{dr}$, for ingoing rays.

The Newman–Penrose transport equations are hierarchical, so we consider them in order. We first focus on the smallest closed subset, which are the Sachs equations (7.100). We make a leading-order ansatz of the form $\rho = \rho_0 r^{n_\rho}$, $\sigma = \sigma_0 r^{n_\sigma}$. Substitution of this ansatz and the leading-order expression for Ψ_0 , given in (7.140), into the Sachs equations yields

$$-n_\rho E \rho_0 r^{n_\rho-1} = \rho_0^2 r^{2n_\rho} + \sigma_0 \bar{\sigma}_0 r^{2n_\sigma}, \quad (7.145)$$

$$-n_\sigma E \sigma_0 r^{n_\sigma-1} = 2\rho_0 \sigma_0 r^{n_\rho+n_\sigma} + 3MKr^{-5}. \quad (7.146)$$

We assume that the term on the right-hand side of (7.146) at $O(r^{n_\rho+n_\sigma})$ is sub-dominant, and can be neglected at leading order. The left-hand side of (7.146) (i.e., the change in σ) is therefore sourced by the curvature term $\Psi_0 = O(r^{-5})$. Balancing these terms forces $n_\sigma = -4$. Rearrangement then fixes the leading-order coefficient $\sigma_0 = \frac{3KM}{4E}$. The shear generated by spacetime curvature through (7.146), in turn, generates expansion ρ through (7.145). Assuming that the term at $O(r^{2n_\rho})$ on the right-hand side of (7.145) is sub-dominant, we find $n_\rho = -7$. Moreover, the coefficient ρ_0 is fixed in terms of σ_0 : $\rho_0 = \frac{1}{7E} \sigma_0 \bar{\sigma}_0$.

Going through the full system of transport equations in this way, we find that the far-field leading-order solutions for the six relevant Newman–Penrose scalars are

$$\rho = \frac{1}{7E} \left(\frac{3KM}{4E} \right)^2 r^{-7} + O(r^{-8}), \quad \sigma = \frac{3KM}{4E} r^{-4} + O(r^{-5}), \quad (7.147)$$

$$\chi = -\frac{iM\sqrt{K}}{\sqrt{2}E} r^{-3} + O(r^{-4}), \quad \tau = -\frac{iM\sqrt{K}}{\sqrt{2}E} r^{-3} + O(r^{-4}), \quad (7.148)$$

$$\mu = -\frac{M}{2E} r^{-2} + O(r^{-3}), \quad \lambda = -\frac{3KM^2}{40E^3} r^{-5} + O(r^{-6}). \quad (7.149)$$

7.6.3 Far-field behaviour of higher-order quantities

We now consider the far-field behaviour of the directional derivatives of Weyl scalars (7.131)–(7.134), which appear on the right-hand side of the transport equations (7.104)–(7.111). Inserting the leading-order solutions for the Newman–Penrose scalars, the Weyl scalars and the null rotation parameter, the leading-order solutions for the directional derivatives of the first two Weyl scalars are

$$\delta\Psi_0 = -\frac{15EM\sqrt{K}a \cos\theta}{\sqrt{2}} r^{-6} + O(r^{-7}), \quad \bar{\delta}\Psi_0 = -6i\sqrt{2}KEMr^{-5} + O(r^{-6}), \quad (7.150)$$

$$\delta\Psi_1 = 6iEMa \cos\theta r^{-5} + O(r^{-6}), \quad \bar{\delta}\Psi_1 = -3EMr^{-4} + O(r^{-5}). \quad (7.151)$$

Now, one may combine the results for the Newman–Penrose scalars, the Weyl scalars and their directional derivatives to calculate leading-order solutions for the higher-order

quantities $\{\delta\mathcal{A}^2, \bar{\delta}\mathcal{A}^2, \delta\rho, \bar{\delta}\rho, \delta\sigma, \bar{\delta}\sigma, \delta\chi, \bar{\delta}\chi\}$. Employing a method similar to that of Section 7.6.2 (i.e., substituting a leading-order ansatz into the transport equations and equating powers of r^{-1}), we find

$$\delta\rho = \frac{9M^2\sqrt{2K^3i}}{56E^2}r^{-7} + O(r^{-8}), \quad \bar{\delta}\rho = -\frac{9M^2\sqrt{2K^3i}}{56E^2}r^{-7} + O(r^{-8}), \quad (7.152)$$

$$\delta\sigma = -\frac{3M\sqrt{2K}a\cos\theta}{2}r^{-5} + O(r^{-6}), \quad \bar{\delta}\sigma = -\frac{3M\sqrt{2K}i}{2}r^{-4} + O(r^{-5}), \quad (7.153)$$

$$\delta\chi = \frac{3Ma\cos\theta}{2}r^{-4} + O(r^{-5}), \quad \bar{\delta}\chi = -Mr^{-3} + O(r^{-4}), \quad (7.154)$$

$$\delta\mathcal{A}^2 = \frac{3M^2\sqrt{2K^3i}}{56E^3}r^{-6} + O(r^{-7}), \quad \bar{\delta}\mathcal{A}^2 = \frac{3M^2\sqrt{2K^3i}}{56E^3}r^{-6} + O(r^{-7}). \quad (7.155)$$

We note that $\bar{\delta}\rho = (\delta\rho)^*$ and $\bar{\delta}\mathcal{A}^2 = (\delta\mathcal{A}^2)^*$, since both ρ and \mathcal{A} are real for a twist-free tetrad.

7.6.4 Far-field behaviour of geometric optics quantities

The geometric optics quantity of interest is $\text{Re}(\mathbf{u}_1)$ which satisfies the transport equation (7.12). The leading-order terms on the right-hand side of this transport equation come from $\bar{\delta}\chi$ and its complex conjugate, which are of order r^{-3} , as given in (7.154). The coefficient of the leading-order term in the far-field expansion for $\bar{\delta}\chi$ is $-M$, which is real. The leading-order terms in the transport equation (7.12) will therefore cancel. In order to determine the leading-order behaviour of $\text{Re}(\mathbf{u}_1)$, it is necessary to consider the sub-leading-order behaviour of the terms on the right-hand side of (7.12).

To find the sub-leading-order part of $\bar{\delta}\chi$ from its transport equation (7.111), we require the sub-leading-order pieces of the terms on the right-hand side. The highest-order terms are the $O(r^{-4})$ and $O(r^{-5})$ pieces of $\bar{\delta}\Psi_1$. Using (7.134), we see that these are, respectively, the $O(r^{-4})$ and $O(r^{-5})$ parts of $(\bar{\delta}C)_{km\bar{m}\bar{m}}$; an expression for the latter quantity is given in (E.65) of Appendix E. Performing a series expansion of this quantity in powers of r^{-1} , we find

$$(\bar{\delta}C)_{km\bar{m}\bar{m}} = -3EMr^{-4} + 24iEMa\cos\theta r^{-5} + O(r^{-6}). \quad (7.156)$$

Using this expression, we find that

$$\bar{\delta}\chi = -Mr^{-3} + 6iMa\cos\theta r^{-4} + O(r^{-5}). \quad (7.157)$$

Crucially, the $O(r^{-4})$ coefficient in (7.157) has a non-zero imaginary part. The term $\bar{\delta}\chi - \delta\bar{\chi}$ will therefore contribute to the right-hand side of the transport equation (7.12) at $O(r^{-4})$. Using this transport equation, we see that the leading-order far-field series solution for $\text{Re}(\mathbf{u}_1)$ is

$$\text{Re}(\mathbf{u}_1) = -\frac{2Ma\cos\theta}{E}r^{-3} + O(r^{-2}). \quad (7.158)$$

We remark that the leading-order piece depends on the spin of the black hole a , and vanishes in the case of a Schwarzschild black hole ($a = 0$).

7.7 Wavefronts, caustic points and transport equations

7.7.1 Wavefronts and caustics in geometric optics

We are interested in the propagation of electromagnetic waves on Kerr spacetime using the (higher-order) geometric optics approximation. Here, we review the important concept of a wavefront in general relativity for a general spacetime (\mathcal{M}, g_{ab}) ; we refer the reader to [7] and references therein for a more detailed overview.

A *wavefront* is a subset of spacetime that can be constructed as follows [7]. First, choose an orientable spacelike two-surface \mathcal{S} . Next, at each point $p \in \mathcal{S}$, choose a null direction orthogonal to \mathcal{S} that depends smoothly on p . Finally, consider all rays which are tangent to the chosen null directions. These null geodesics are the *generators* of the wavefront. The wavefront \mathcal{W} is the union of all generators. If \mathcal{S} is a two-sphere of radius ε centred on $p \in \mathcal{M}$, then in the limit $\varepsilon \rightarrow 0$, the wavefront is the light cone with vertex p .

One may describe a wavefront using a coordinate system which is well-adapted to the behaviour of the generators [172]. We allow the affine parameter s to be one of the coordinates, and we introduce two additional (local) coordinates $\{\theta^1, \theta^2\}$ to label the null generators; the latter are constant along each generator, and span the two-surface \mathcal{S} (which is transverse to the generators). Each null generator $\gamma(s)$ is chosen to be affinely parametrised such that $\gamma(s_0) \in \mathcal{S}$ and the tangent vector field k^a depends smoothly on the point $\gamma(s_0)$. The wavefront \mathcal{W} is then the image of a map

$$(s, \theta^1, \theta^2) \mapsto W^a(s, \theta^1, \theta^2), \quad a \in \{0, 1, 2, 3\}. \quad (7.159)$$

In the neighbourhood of the two-surface \mathcal{S} (i.e., close to $s = s_0$), the map (7.159) is an embedding, so \mathcal{W} is a three-dimensional (null) submanifold of \mathcal{M} .² However, away from \mathcal{S} , neighbouring null generators may intersect each other and the wavefront may fail to be a submanifold.

The *caustic* of the wavefront is defined as the set of all points where (7.159) fails to be an immersion; in other words, the rank of the differential of the map (7.159) is not maximal (i.e., it is strictly less than 3) [7]. Heuristically, a caustic is a transverse self-intersection of the wavefront. We note that the derivative of the map (7.159) with respect to the affine parameter s is always non-zero, so the rank of the map cannot be zero. If the rank is $3 - 1 = 2$, then the caustic point is said to have *multiplicity one*. If, on the other hand, the rank is $3 - 2 = 1$, the caustic point is said to be of *multiplicity two*. In the former case, the cross-section of an infinitesimally thin bundle of null rays (generators) in the instantaneous wavefront collapses down to a line at a caustic point. In the latter case, the bundle of generators collapses down to a single point.

If the wavefront is a light cone with vertex p , caustic points are said to be *conjugate* to the point p along a generator (i.e., a null geodesic) [7]. For an arbitrary wavefront, a caustic point is said to be *conjugate* to a spacelike two-surface in the initial wavefront \mathcal{S} . The terms *conjugate point* and *caustic point* are therefore used interchangeably.

²A submanifold is said to be *null* (or *lightlike*) if its induced metric is degenerate. Since k^a is orthogonal to \mathcal{S} , it follows that the induced metric on the wavefront close to \mathcal{S} is degenerate [172]. It is therefore a null submanifold.

Away from caustic points, a wavefront is a hypersurface (i.e., a three-dimensional sub-manifold), which can be expressed as a function of the spacetime coordinates in the form $S(x^a) = \text{constant}$, where S satisfies the *Hamilton–Jacobi equation* (or *eikonal equation*) $g^{ab}S_{,a}S_{,b} = 0$. However, at caustic points, a wavefront will typically exhibit complicated structure, with “cuspidal edges”, or “vertices”. Caustics can be classified locally using the singularity theory of Arnold [84, 316].

Recall from Section 3.3.2 that the lensing dynamics of an infinitesimal bundle of rays in a wavefront \mathcal{S} can be understood by analysing the geodesic deviation equation (or Jacobi equation), given in (3.36). The cross-section of a bundle of rays is described by the shape parameters d_{\pm} . For an infinitesimal twist-free bundle of generators, a point where only one of the two shape parameters vanishes is a caustic point of multiplicity one; a point where both of the shape parameters vanish is a caustic point of multiplicity two [7]. It is clear from this definition that, at a caustic point, the cross-sectional area of the bundle $A = \pi d_+ d_-$ vanishes.

The behaviour of the optical scalars in the neighbourhood of caustic points can be determined by considering the transport equations (7.100), i.e., the Sachs equations. Suppose there is a caustic point at $s = s_0$. For a caustic point of multiplicity one, the optical scalars are [7, 317]

$$\rho(s) = -\frac{1}{2}(s - s_0)^{-1} + O(s - s_0), \quad |\sigma(s)| = \frac{1}{2}(s - s_0)^{-1} + O(s - s_0). \quad (7.160)$$

For a caustic point of multiplicity two, the optical scalars are given by [7, 317]

$$\rho(s) = (s - s_0)^{-1} + O(s - s_0), \quad \sigma(s) = \frac{1}{3}\Psi_0(s_0)(s - s_0) + O((s - s_0)^2). \quad (7.161)$$

Caustic points of multiplicity one are somewhat more generic than caustic points of multiplicity two. Hereafter, when we use the term caustic point, we mean a caustic point of multiplicity one, unless stated otherwise.

7.7.2 Behaviour of Newman–Penrose scalars at caustic points

It is clear from (7.160) that the optical scalars diverge like $(s - s_0)^{-1}$ at a caustic point. In the vicinity of a caustic point, the behaviour of the full set of complex Newman–Penrose scalars can be deduced from the system of transport equations (7.100)–(7.102) as follows.

In particular, let us assume that we have a caustic point of multiplicity one at $s = 0$, where s is the affine parameter. (The affine parameter can always be redefined such that the caustic point occurs at $s = 0$.) We suppose that each Newman–Penrose scalar $z \in \{\rho, \sigma, \chi, \tau, \mu, \lambda\}$ can be expanded as a power series of the form

$$z(s) = s^{n_z} \sum_{j=0}^{\infty} z_j s^j, \quad (7.162)$$

where n_z is some exponent to be determined for each Newman–Penrose scalar z , and the coefficients z_j are complex. Moreover, we expand the Weyl scalars, which appear on the right-hand side of the Newman–Penrose transport equations (7.100)–(7.102), as a power

series of the form

$$\Psi_i(s) = \sum_{j=0}^{\infty} \psi_i^{(j)} s^j, \quad (7.163)$$

where $\psi_i^{(j)}$ are complex coefficients. (We note that the Weyl scalars are regular at caustic points.)

We insert the ansatz (7.162) and (7.163) into the hierarchical system of transport equations (7.100)–(7.102) and consider each closed subsystem in turn (see the method of Section 7.6.2). Balancing the leading order terms on both sides of the transport equations fixes the exponents $n_z = -1$ for all $z \in \{\rho, \sigma, \chi, \tau, \lambda, \mu\}$. We find that, close to a caustic point of multiplicity one at $s = 0$, the behaviour of the Newman–Penrose scalars through sub-leading order is given by

$$\rho(s) = -\frac{1}{2}s^{-1} + \rho_1 + O(s), \quad \sigma(s) = -\frac{1}{2}e^{2i\varphi}s^{-1} - e^{2i\varphi}\rho_1 + O(s), \quad (7.164)$$

$$\chi(s) = |\chi_0|ie^{i\varphi}s^{-1} + |\chi_1|e^{i\varphi} + O(s), \quad \tau(s) = |\tau_0|e^{i\varphi}s^{-1} + |\tau_1|ie^{i\varphi} + O(s), \quad (7.165)$$

$$\mu(s) = \mu_0s^{-1} + \mu_1 + O(s), \quad \lambda(s) = \mu_0e^{-2i\varphi}s^{-1} - \mu_1e^{-2i\varphi} + O(s), \quad (7.166)$$

where $\varphi \in [0, 2\pi)$, $\rho_1, |\chi_0|, |\chi_1|, |\tau_0|, |\tau_1| \in \mathbb{R}$ and $\mu_0, \mu_1 \in \mathbb{C}$ are free parameters determined by the ray along which the transport equations (7.100)–(7.102) are evolved. The Weyl scalars (7.163), which are $O(1)$ quantities, do not contribute to the series solutions (7.164)–(7.166) until $O(s)$, i.e., sub-sub-leading order. Hence, close to a caustic point of multiplicity one, spacetime curvature does not play a significant role in the behaviour of the Newman–Penrose scalars. The leading-order results for the optical scalars, given in (7.164), are equivalent to those presented by Perlick [7] and Seitz *et al.* [317], cf. (7.160).

The cross-sectional area A and geometric optics square-amplitude \mathcal{A}^2 satisfy the transport equations $DA = -2\rho A$ and $D\mathcal{A}^2 = 2\rho\mathcal{A}^2$. Close to the caustic point, we expand these quantities as generalised power series of the form

$$A(s) = \sum_{j=0}^{\infty} A_j s^j, \quad \mathcal{A}^2(s) = s^n \sum_{j=0}^{\infty} a_j s^j, \quad (7.167)$$

where n is some index to be determined. Inserting (7.167) into the above transport equations, and noting that the cross-sectional area vanishes at the caustic point (i.e., $A(0) = 0$), we find the sub-leading-order solutions

$$A(s) = A_1 s - 2\rho_1 A_1 s^2 + O(s^3), \quad (7.168)$$

$$\mathcal{A}^2(s) = a_0 s^{-1} + 2\rho_1 a_0 + O(s), \quad (7.169)$$

where $A_1, a_0 \in \mathbb{R}$ are free parameters. The square-amplitude \mathcal{A}^2 therefore diverges like $O(s^{-1})$ close to a caustic point of multiplicity one. We use the square of the amplitude (rather than the amplitude \mathcal{A}) to avoid fractional powers in the generalised power series expansions close to the caustic point.

7.7.3 Behaviour of higher-order quantities at caustic points

We now turn our attention to the quantities $\{\delta\rho, \bar{\delta}\rho, \delta\sigma, \bar{\delta}\sigma, \delta\chi, \bar{\delta}\chi\}$, which are governed by the transport equations (7.106)–(7.111). At first glance, this system of equations appears to be non-linear. However, we note that, if the Newman–Penrose scalars are known,

then the transport equations (7.106)–(7.111) may be coupled to their complex conjugates to yield a linear system of first-order ordinary differential equations for the unknown quantities and their complex conjugates. This system may be expressed as a linear inhomogeneous system of ordinary differential equations of the form

$$D\mathbf{x} = M\mathbf{x} + \mathbf{S} + \mathbf{W}, \quad (7.170)$$

where

$$\mathbf{x} = [\delta\rho, (\delta\rho)^*, \delta\sigma, (\delta\sigma)^*, \bar{\delta}\sigma, (\bar{\delta}\sigma)^*, \delta\chi, (\delta\chi)^*, \bar{\delta}\chi, (\bar{\delta}\chi)^*], \quad (7.171)$$

is a vector which contains all of the dependent variables and their complex conjugates. We recall that $(\delta\rho)^* = \bar{\delta}\rho$, since ρ is real. In (7.170), M is a coefficient matrix of Newman–Penrose scalars, which is given by

$$M = \begin{bmatrix} 3\rho & \sigma & \bar{\sigma} & 0 & 0 & \sigma & 0 & 0 & 0 & 0 \\ \bar{\sigma} & 3\rho & 0 & \sigma & \bar{\sigma} & 0 & 0 & 0 & 0 & 0 \\ 2\sigma & 0 & 3\rho & 0 & \sigma & 0 & 0 & 0 & 0 & 0 \\ 0 & 2\bar{\sigma} & 0 & 3\rho & 0 & \bar{\sigma} & 0 & 0 & 0 & 0 \\ 0 & 2\sigma & \bar{\sigma} & 0 & 3\rho & 0 & 0 & 0 & 0 & 0 \\ 2\bar{\sigma} & 0 & 0 & \sigma & 0 & 3\rho & 0 & 0 & 0 & 0 \\ \chi & 0 & -\bar{\chi} & 0 & 0 & 0 & 2\rho & 0 & \sigma & -\sigma \\ 0 & \bar{\chi} & 0 & -\chi & 0 & 0 & 0 & 2\rho & -\bar{\sigma} & \bar{\sigma} \\ 0 & \chi & 0 & 0 & -\bar{\chi} & 0 & \bar{\sigma} & -\sigma & 2\rho & 0 \\ \bar{\chi} & 0 & 0 & 0 & 0 & -\chi & -\bar{\sigma} & \sigma & 0 & 2\rho \end{bmatrix}. \quad (7.172)$$

The components of the inhomogeneous terms which appear on the right-hand side of (7.170) are given by

$$\mathbf{S} = [(\rho^2 + \sigma\bar{\sigma})\tau, (*), 2\rho\sigma\tau, (*), 2\rho\sigma\bar{\tau}, (*), \tau(\rho\chi - \sigma\bar{\chi}), (*), \bar{\tau}(\rho\chi - \sigma\bar{\chi}), (*)], \quad (7.173)$$

$$\mathbf{W} = [0, (*), \delta\Psi_0 - \tau\Psi_0, (*), \bar{\delta}\Psi_0 - \bar{\tau}\Psi_0, (*), \delta\Psi_1 - \tau\Psi_1, (*), \bar{\delta}\Psi_1 - \bar{\tau}\Psi_1, (*)], \quad (7.174)$$

where an asterisk in parentheses (*) is used for brevity to denote the complex conjugate of the preceding component of the vector. For example, the second component of \mathbf{S} is $((\rho^2 + \sigma\bar{\sigma})\tau)^* = (\rho^2 + \sigma\bar{\sigma})\bar{\tau}$. The vector \mathbf{S} contains all of the Newman–Penrose scalars, and \mathbf{W} contains all of the curvature terms, i.e., the Weyl curvature scalars and their directional derivatives.

Our aim is to find a generalised power series solution for the higher-order Newman–Penrose quantities $\{\delta\rho, \bar{\delta}\rho, \delta\sigma, \bar{\delta}\sigma, \delta\chi, \bar{\delta}\chi\}$ through sub-leading order, in the vicinity of a caustic point of multiplicity one. To do this, we seek a generalised power series solution to the linear inhomogeneous system (7.170), with the additional constraint that the solution must satisfy the complex conjugacy property $x_{2k-1} = \bar{x}_{2k}$ for all $k \in \{1, 2, 3, 4, 5\}$, where x_k are the coefficients of the solution \mathbf{x} . We expand \mathbf{x} as a generalised power series of the form

$$\mathbf{x}(s) = s^n \sum_{j=0}^{\infty} \mathbf{x}_j s^j, \quad (7.175)$$

where \mathbf{x}_j are vectors of complex coefficients and the index n is to be determined from analysis of the system (7.170). We note here that since the solution \mathbf{x} must satisfy the

complex conjugacy property (by construction), each of the coefficients \mathbf{x}_j at $O(s^{j+n})$ must also satisfy the same property. Upon substitution of (7.164)–(7.166) into (7.172)–(7.174), we see that $M = O(s^{-1})$, $\mathbf{S} = O(s^{-3})$ and $\mathbf{W} = O(s^{-1})$. We therefore write

$$M(s) = s^{-1} \sum_{j=0}^{\infty} M_j s^j, \quad \mathbf{S}(s) = s^{-3} \sum_{j=0}^{\infty} \mathbf{S}_j s^j, \quad \mathbf{W}(s) = s^{-3} \sum_{j=0}^{\infty} \mathbf{W}_j s^j. \quad (7.176)$$

Here, M_j are matrices given by (7.172) with the Newman–Penrose scalars replaced by their $O(s^j)$ coefficients from the series expansions (7.162). Similarly, \mathbf{S}_j and \mathbf{W}_j are vectors given by (7.173) and (7.174), respectively, with the Newman–Penrose scalars and curvature terms replaced by their $O(s^j)$ coefficients.

At leading order, the system of equations (7.170) is

$$M_0 \mathbf{x}_0 = n \mathbf{x}_0, \quad (7.177)$$

where we have assumed that the $O(s^{-3})$ terms are sub-dominant. This is an eigenvalue problem for the matrix M_0 . The characteristic polynomial for the matrix M_0 is

$$n^2 (n+1)^4 (n+2)^3 (n+3) = 0, \quad (7.178)$$

from which the eigenvalues and their multiplicities can be read off straightforwardly. The smallest eigenvalue of M_0 is $n = -3$, so $\mathbf{x} = O(s^{-3})$. The eigenvector of the matrix M_0 corresponding to the eigenvalue $n = -3$ is

$$\mathbf{v}^{(-3)} = [e^{i\varphi}, (*), e^{3i\varphi}, (*), e^{i\varphi}, (*), -2|\chi_0|ie^{2i\varphi}, (*), -2i|\chi_0|, (*)], \quad (7.179)$$

where we have again used the notation $(*)$ to denote the complex conjugate of the previous term. The eigenvector (7.179) clearly satisfies the complex conjugacy property. The leading-order solution is therefore any (complex) scalar multiple of this eigenvector: $\mathbf{x}_0 = \alpha_0 \mathbf{v}^{(-3)}$. Demanding that the leading-order solution \mathbf{x}_0 satisfies the complex conjugacy property means that α_0 must be a real quantity. To leading order, the solutions to the transport equations (7.106)–(7.111) close to a caustic point of multiplicity one are therefore given by

$$\delta\rho(s) = \alpha_0 e^{i\varphi} s^{-3} + O(s^{-2}), \quad \bar{\delta}\rho(s) = \alpha_0 e^{i\varphi} s^{-3} + O(s^{-2}), \quad (7.180)$$

$$\delta\sigma(s) = \alpha_0 e^{3i\varphi} s^{-3} + O(s^{-2}), \quad \bar{\delta}\sigma(s) = \alpha_0 e^{i\varphi} s^{-3} + O(s^{-2}), \quad (7.181)$$

$$\delta\chi(s) = -2\alpha_0 |\chi_0| i e^{2i\varphi} s^{-3} + O(s^{-2}), \quad \bar{\delta}\chi(s) = -2\alpha_0 |\chi_0| i s^{-3} + O(s^{-2}). \quad (7.182)$$

The sub-leading-order terms in (7.170) yield an inhomogeneous system,

$$(M_0 + 2I) \mathbf{x}_1 = -(\mathbf{S}_0 + M_1 \mathbf{x}_0), \quad (7.183)$$

where I denotes the identity matrix. Interestingly, we find that $\mathbf{x}_0 \in \text{Null}(M_1)$, where $\text{Null}(M_1)$ denotes the null space of the matrix M_1 . Hence, the term $M_1 \mathbf{x}_0$ on the right-hand side of (7.183) vanishes. The matrix $M_0 + 2I$ is singular since $n = -2$ is an eigenvalue of M_0 . One can see from (7.178) that this eigenvalue has multiplicity three. Solving the system (7.183) using Gaussian elimination, we arrive at the sub-leading-order solution

$$\mathbf{x}_1 = \sum_{k=1}^3 w_k \mathbf{v}_k^{(-2)} + \mathbf{y}_1, \quad (7.184)$$

where $w_k = \xi_k + i\eta_k$ are arbitrary complex coefficients. The three eigenvectors of M_0 corresponding to the eigenvalue $n = -2$ are

$$\mathbf{v}_1^{(-2)} = [ie^{i\varphi}, -ie^{-i\varphi}, 3ie^{3i\varphi}, -3ie^{-3i\varphi}, ie^{i\varphi}, -ie^{-i\varphi}, 4|\chi_0|e^{2i\varphi}, 4|\chi_0|e^{-2i\varphi}, 0, 0], \quad (7.185)$$

$$\mathbf{v}_2^{(-2)} = [0, 0, ie^{3i\varphi}, -ie^{-3i\varphi}, ie^{i\varphi}, -ie^{-i\varphi}, 2|\chi_0|e^{2i\varphi}, 0, 2|\chi_0|, 0], \quad (7.186)$$

$$\mathbf{v}_3^{(-2)} = [ie^{i\varphi}, -ie^{-i\varphi}, ie^{3i\varphi}, -ie^{-3i\varphi}, -ie^{i\varphi}, ie^{-i\varphi}, 4|\chi_0|e^{2i\varphi}, 0, 0, -4|\chi_0|], \quad (7.187)$$

which are clearly linearly independent. The particular solution to the inhomogeneous system (7.183) is given by

$$\mathbf{y}_1 = -\frac{1}{2}|\tau_0| [e^{i\varphi}, e^{-i\varphi}, e^{3i\varphi}, e^{-3i\varphi}, e^{i\varphi}, e^{-i\varphi}, 0, 0, 0, 0]. \quad (7.188)$$

We note here that the particular solution (7.188) satisfies the complex conjugacy property; however, the eigenvectors (7.186) and (7.187) do not. Demanding that the components of the sub-leading-order solution (7.184) satisfy the complex conjugacy property gives a relationship between the complex coefficients w_k :

$$w_1 = \xi_1 + i\eta_1, \quad w_2 = \xi_2 - i\eta_1, \quad w_3 = -2\xi_2 - 2i\eta_1, \quad (7.189)$$

where $\xi_1, \xi_2, \eta_1 \in \mathbb{R}$ are free parameters.

The generalised power series solutions for the higher-order Newman–Penrose quantities through sub-leading order are given by

$$\delta\rho(s) = \alpha_0 e^{i\varphi} s^{-3} + \left(\xi_1 + \xi_2 + \frac{i}{2}|\tau_0| \right) ie^{i\varphi} s^{-2} + O(s^{-1}), \quad (7.190)$$

$$\bar{\delta}\rho(s) = \alpha_0 e^{i\varphi} s^{-3} - \left(\xi_1 + \xi_2 - \frac{i}{2}|\tau_0| \right) ie^{-i\varphi} s^{-2} + O(s^{-1}), \quad (7.191)$$

$$\delta\sigma(s) = \alpha_0 e^{3i\varphi} s^{-3} + \left(3\xi_1 - \xi_2 - \frac{i}{2}|\tau_0| \right) ie^{3i\varphi} s^{-2} + O(s^{-1}), \quad (7.192)$$

$$\bar{\delta}\sigma(s) = \alpha_0 e^{i\varphi} s^{-3} + \left(\xi_1 - 3\xi_2 - \frac{i}{2}|\tau_0| \right) ie^{i\varphi} s^{-2} + O(s^{-1}), \quad (7.193)$$

$$\delta\chi(s) = -2\alpha_0 |\chi_0| ie^{2i\varphi} s^{-3} + 4|\chi_0| (\xi_1 - i\eta_1) e^{2i\varphi} s^{-2} + O(s^{-1}), \quad (7.194)$$

$$\bar{\delta}\chi(s) = -2\alpha_0 |\chi_0| is^{-3} - 4|\chi_0| (\xi_2 + i\eta_1) s^{-2} + O(s^{-1}). \quad (7.195)$$

These solutions feature three free parameters ($\varphi, |\chi_0|, |\tau_0| \in \mathbb{R}$) which appear in the leading-order solutions for the Newman–Penrose scalars; one free parameter ($\alpha_0 \in \mathbb{R}$) which comes from the leading-order part; and three free parameters ($\xi_1, \xi_2, \eta_1 \in \mathbb{R}$) which arise from the sub-leading-order part.

As described in Section 7.5.4, we have the pair of algebraic constraints $Z = 0$ and $W = 0$, given in (7.135) and (7.135), which must be satisfied by the Newman–Penrose scalars, the higher-order Newman–Penrose quantities, and the Weyl scalars. We now impose these constraints on the solutions (7.190)–(7.195) through $O(s^{-2})$. The dominant terms in (7.135) are the $O(s^{-3})$ pieces of $\delta\rho$ and $\bar{\delta}\sigma$. Remarkably, these are equal, so the constraint $Z = 0$ is satisfied at leading order. Similarly, the dominant terms in (7.136) are the $O(s^{-3})$ pieces of $\bar{\delta}\chi$ and its complex conjugate. Since these quantities are purely

imaginary, their sum vanishes, and the constraint $W = 0$ is satisfied at leading order. In order for the constraints (7.135) and (7.136) to be satisfied at $O(s^{-2})$, we require

$$\xi_2 = \frac{1}{4} |\chi_0|. \quad (7.196)$$

Both $Z = 0$ and $W = 0$ yield the same condition. Replacing the parameter ξ_2 in (7.190)–(7.195) using (7.196) yields the generalised power series solutions through sub-leading order for the quantities $\{\delta\rho, \bar{\delta}\rho, \delta\sigma, \bar{\delta}\sigma, \delta\chi, \bar{\delta}\chi\}$.

Using the sub-leading-order solutions for the Newman–Penrose scalars and the higher-order Newman–Penrose quantities, we may also look for a generalised series solution to the transport equations (7.104) and (7.105) close to a caustic point of multiplicity one. Noting that $\bar{\delta}\mathcal{A}^2 = (\delta\mathcal{A}^2)^*$, we seek a solution of the form

$$\delta\mathcal{A}^2(s) = s^n \sum_{j=0}^{\infty} \mathbf{a}_j s^j, \quad (7.197)$$

where \mathbf{a}_j are complex coefficients, and the index n is to be determined by balancing terms in (7.104) and (7.105). Inserting (7.197) and its complex conjugate into the transport equations (7.104) and (7.105), we find $n = -3$; hence, $\delta\mathcal{A}^2$ diverges like s^{-3} at a caustic point of multiplicity one.

At $O(s^{-3})$, the equations (7.104) and (7.105) lead to

$$\begin{bmatrix} 3\rho_0 + 3 & \sigma_0 \\ \bar{\sigma}_0 & 3\rho_0 + 3 \end{bmatrix} \begin{bmatrix} \mathbf{a}_0 \\ \bar{\mathbf{a}}_0 \end{bmatrix} = -2a_0\alpha_0 \begin{bmatrix} e^{i\varphi} \\ e^{-i\varphi} \end{bmatrix}. \quad (7.198)$$

The matrix on the left-hand side of this system is invertible, and it follows that $\mathbf{a}_0 = -2a_0\alpha_0 e^{i\varphi}$. Hence $\delta\mathcal{A}^2 = -2a_0\alpha_0 e^{i\varphi} s^{-3} + O(s^{-2})$. At $O(s^{-2})$, the equations (7.104) and (7.105) yield

$$\begin{bmatrix} 3\rho_0 + 2 & \sigma_0 \\ \bar{\sigma}_0 & 3\rho_0 + 2 \end{bmatrix} \begin{bmatrix} \mathbf{a}_1 \\ \bar{\mathbf{a}}_1 \end{bmatrix} = \frac{1}{2} a_0 (4\xi_1 + |\chi_0|) \begin{bmatrix} -ie^{i\varphi} \\ ie^{-i\varphi} \end{bmatrix}. \quad (7.199)$$

This system is non-invertible; however, it may be solved using Gaussian elimination, for example. The solution is $\mathbf{a}_1 = [\zeta_1 + \frac{ia_0}{2} (4\xi_1 + |\chi_0|)] e^{i\varphi}$, where $\zeta_1 \in \mathbb{R}$ is a free parameter. The general solution to (7.104) and (7.105) through sub-leading order is therefore

$$\delta\mathcal{A}^2(s) = -2a_0\alpha_0 e^{i\varphi} s^{-3} + \left[\zeta_1 + \frac{ia_0}{2} (4\xi_1 + |\chi_0|) \right] e^{i\varphi} s^{-2} + O(s^{-1}), \quad (7.200)$$

with $\bar{\delta}\mathcal{A}^2$ given by the complex conjugate of (7.200).

Inserting all of the relevant sub-leading-order series solutions into the right-hand side of the transport equation (7.12), we find

$$D(\text{Re}(\mathbf{u}_1)) = 4\alpha_0 |\chi_0| s^{-3} - \left[4|\chi_0| (\alpha_0\rho_1 - \eta_1) + \text{Im}(\mu_0) + \frac{\zeta_1 |\chi_0|}{a_0} \right] s^{-2} + O(s^{-1}). \quad (7.201)$$

Integration of this transport equation with respect to s gives

$$\text{Re}(\mathbf{u}_1(s)) = -2\alpha_0 |\chi_0| s^{-2} + \left[4|\chi_0| (\alpha_0\rho_1 - \eta_1) + \text{Im}(\mu_0) + \frac{\zeta_1 |\chi_0|}{a_0} \right] s^{-1} + O(1). \quad (7.202)$$

Hence, the quantity $\text{Re}(\mathbf{u}_1)$ diverges like s^{-2} at a caustic point of multiplicity one. The $O(s^{-2})$ part of the divergence is controlled by the parameters $|\chi_0|$ and α_0 ; the former arises in the leading-order solution for the Newman–Penrose scalars (7.164)–(7.166), whilst the latter comes from the leading-order series solution for the higher-order Newman–Penrose quantities (7.180)–(7.182). The $O(s^{-1})$ piece features the parameters ρ_1 , $|\chi_0|$ and μ_0 , which come from the leading-order solutions (7.164)–(7.166); a_0 , which comes from the leading-order part of (7.169); α_0 and η_1 , which arise from the leading- and sub-leading-order solutions (7.190)–(7.195), respectively; and ζ_1 , which appears in the sub-leading-order piece of (7.200).

We remark here that a quantity with spin weight \mathfrak{s} admits a power series expansion close to the caustic point whose coefficients are proportional to the unit complex number $e^{i\mathfrak{s}\pi}$. (We refer the reader to Section 2.5.6 for a brief review of the Geroch–Held–Penrose calculus and the notion of spin weight.)

The transport equations for which we have obtained series solutions come from the higher-order geometric optics formalism of Dolan [81], which assumes only that the spacetime is Ricci-flat. The results obtained in Section 7.7.2 and the present section are general, in that they apply to any Ricci-flat manifold. The curvature terms (i.e., Weyl scalars and their derivatives) do not play a role until sub-sub-leading order.

7.8 Evolving transport equations through caustic points

Our primary aim in this chapter is to calculate higher-order geometrical optics corrections to the stress–energy tensor (see Section 7.2.2). This involves calculating the quantity $\text{Re}(\mathbf{u}_1)$, for which we have a transport equation involving Newman–Penrose scalars and higher-order Newman–Penrose quantities (e.g. $\bar{\delta}\chi$). When evolving the relevant system of transport equations for an infinitesimal bundle of rays, we will generically encounter caustic points, where neighbouring rays cross. As demonstrated in Section 7.7, the Newman–Penrose scalars, higher-order Newman–Penrose quantities and geometric optics quantities diverge at caustic points; we therefore require a method which allows us to evolve the transport equations through caustic points, so that the quantities of interest, such as $\text{Re}(\mathbf{u}_1)$, can be calculated after scattering has occurred (at $r \rightarrow \infty$).

As reviewed in Section 3.3.2, Dolan [81, 171] demonstrated that certain Newman–Penrose scalars – namely $\{\rho, \sigma, \chi, \tau, \mu, \lambda\}$ – can be calculated by decomposing the geodesic deviation equation (3.36) and the differential precession equation (3.49) in a parallel-transported null tetrad $\{k, n, m, \bar{m}\}$. This approach yields a system of first- and second-order ordinary differential equations which *can* be evolved through caustic points. The Newman–Penrose scalars can then be found by inverting a system of simultaneous equations for a pair of linearly independent solutions to the aforementioned differential equations. This inversion breaks down at caustic points; however, it is well-defined everywhere else along the central ray of the bundle. The Newman–Penrose scalars may therefore be calculated *after* the first caustic point using this method.

Whilst it is possible that a similar formalism could be used to find *all* of the Newman–Penrose and geometrical optics quantities along an entire ray (i.e., past the first caustic

point), it is not immediately obvious which equations (cf. geodesic deviation and differential precession) should be considered for this purpose. We therefore proceed by proposing a more pragmatic approach to evolving the full system of transport equations through caustic points, which relies on “regularising” the divergent quantities. In Section 7.8.1, we introduce the regularisation method, which could be applied to any parallel-propagated twist-free complex null tetrad on a Ricci-flat spacetime. In Section 7.8.2, we apply the numerical method to Kerr spacetime, and present some preliminary results. We comment on these results in Section 7.9.

7.8.1 Regularisation of transport equations

In this section, we assume that there is a caustic point (of multiplicity one) at $s = s_0$, where s is the affine parameter along the central ray. We note first that, close to a caustic point, the cross-sectional area scales like $(s - s_0)$; see (7.168).

Consider the Newman–Penrose scalars $\{\rho, \sigma, \chi, \tau, \mu, \lambda\}$, which obey the transport equations (7.100)–(7.102), and which admit generalised power series solutions in the neighbourhood of the caustic point given by (7.164)–(7.166) through sub-leading order. Each Newman–Penrose scalar diverges like $(s - s_0)^{-1}$ as we approach the caustic point. We choose to regularise the Newman–Penrose scalars by multiplying each one by the cross-sectional area, e.g. $\hat{\rho} = A\rho$, to cancel out the $O((s - s_0)^{-1})$ divergence. The regularised Newman–Penrose scalars, which do not diverge at caustic points, satisfy the system of transport equations

$$D\hat{\rho} = \frac{1}{A} \left(\hat{\sigma}\bar{\hat{\sigma}} - \hat{\rho}^2 \right), \quad D\hat{\sigma} = A\Psi_0, \quad (7.203)$$

$$D\hat{\chi} = -\frac{1}{A} \left(\hat{\rho}\hat{\chi} + \hat{\sigma}\bar{\hat{\chi}} \right) + A\Psi_1, \quad D\hat{\tau} = \frac{1}{A} \left(\hat{\sigma}\bar{\hat{\tau}} - \hat{\rho}\hat{\tau} \right) + A\Psi_1, \quad (7.204)$$

$$D\hat{\lambda} = \frac{1}{A} \left(\hat{\sigma}\hat{\mu} - \hat{\rho}\hat{\lambda} \right), \quad D\hat{\mu} = \frac{1}{A} \left(\hat{\sigma}\hat{\lambda} - \hat{\rho}\hat{\mu} \right) + A\Psi_2. \quad (7.205)$$

We note that the cross-sectional area now satisfies the transport equation

$$DA = -2\hat{\rho}. \quad (7.206)$$

As in the case of the Newman–Penrose scalars, the square-amplitude \mathcal{A}^2 diverges like $(s - s_0)^{-1}$, so we regularise via $\hat{\mathcal{A}}^2 = A\mathcal{A}^2$. The regularised square-amplitude then satisfies

$$D\hat{\mathcal{A}}^2 = 0, \quad (7.207)$$

i.e., it is constant. We remark that the notation used here is somewhat misleading: \mathcal{A}^2 is not necessarily a non-negative quantity. One can see from the series solution (7.168) that A will change sign at a caustic point; since \mathcal{A}^2 is proportional to A^{-1} , this quantity will also change sign at a caustic point.

Now consider the higher-order quantities $\{\delta\rho, \bar{\delta}\rho, \delta\sigma, \bar{\delta}\sigma, \delta\chi, \bar{\delta}\chi, \delta\mathcal{A}^2, \bar{\delta}\mathcal{A}^2\}$. The generalised series solutions for these quantities are given in (7.190)–(7.195) and (7.200). Each of these quantities diverges like $(s - s_0)^{-3}$ at a caustic point. We may therefore regularise these quantities by multiplying by A^3 , e.g. $\hat{\delta\rho} = A^3\delta\rho$. The rescaled quantities satisfy the

system of equations

$$D(\widehat{\delta\rho}) = \frac{1}{A} \left(-3\widehat{\rho}\widehat{\delta\rho} + \widehat{\sigma}\widehat{\delta\rho} + \overline{\widehat{\sigma}}\widehat{\delta\sigma} + \widehat{\sigma}(\widehat{\delta\sigma})^* \right) - \left(\widehat{\rho}^2 + \widehat{\sigma}\overline{\widehat{\sigma}} \right) \widehat{\tau}, \quad (7.208)$$

$$D(\widehat{\delta\rho}) = \frac{1}{A} \left(-3\widehat{\rho}\widehat{\delta\rho} + \overline{\widehat{\sigma}}\widehat{\delta\rho} + \widehat{\sigma}\widehat{\delta\sigma} + \widehat{\sigma}(\widehat{\delta\sigma})^* \right) - \left(\widehat{\rho}^2 + \widehat{\sigma}\overline{\widehat{\sigma}} \right) \overline{\widehat{\tau}}, \quad (7.209)$$

$$D(\widehat{\delta\sigma}) = \frac{1}{A} \left(-3\widehat{\rho}\widehat{\delta\sigma} + 2\widehat{\sigma}\widehat{\delta\rho} + \widehat{\sigma}\widehat{\delta\sigma} \right) - 2\widehat{\rho}\widehat{\sigma}\widehat{\tau} - A^2\widehat{\tau}\Psi_0 + A^3\delta\Psi_0, \quad (7.210)$$

$$D(\widehat{\delta\sigma}) = \frac{1}{A} \left(-3\widehat{\rho}\widehat{\delta\sigma} + 2\widehat{\sigma}\widehat{\delta\rho} + \overline{\widehat{\sigma}}\widehat{\delta\sigma} \right) - 2\widehat{\rho}\widehat{\sigma}\overline{\widehat{\tau}} - A^2\overline{\widehat{\tau}}\Psi_0 + A^3\delta\Psi_0, \quad (7.211)$$

$$D(\widehat{\delta\chi}) = \frac{1}{A} \left(-4\widehat{\rho}\widehat{\delta\chi} + \widehat{\chi}\widehat{\delta\rho} + \widehat{\sigma}\widehat{\delta\chi} - \widehat{\sigma}(\widehat{\delta\chi})^* - \overline{\widehat{\chi}}\widehat{\delta\sigma} \right) - \left(\widehat{\rho}\widehat{\chi} - \widehat{\sigma}\overline{\widehat{\chi}} \right) \widehat{\tau} \\ - A^2\widehat{\tau}\Psi_1 + A^3\delta\Psi_1, \quad (7.212)$$

$$D(\widehat{\delta\chi}) = \frac{1}{A} \left(-4\widehat{\rho}\widehat{\delta\chi} + \widehat{\chi}\widehat{\delta\rho} + \overline{\widehat{\sigma}}\widehat{\delta\chi} - \widehat{\sigma}(\widehat{\delta\chi})^* - \overline{\widehat{\chi}}\widehat{\delta\sigma} \right) - \left(\widehat{\rho}\widehat{\chi} - \widehat{\sigma}\overline{\widehat{\chi}} \right) \overline{\widehat{\tau}} \\ - A^2\overline{\widehat{\tau}}\Psi_1 + A^3\delta\Psi_1, \quad (7.213)$$

$$D(\widehat{\delta\mathcal{A}^2}) = \frac{1}{A} \left(2\widehat{\mathcal{A}^2}\widehat{\delta\rho} - 3\widehat{\rho}\widehat{\delta\mathcal{A}^2} + \widehat{\sigma}\widehat{\delta\mathcal{A}^2} \right) - 2\widehat{\mathcal{A}^2}\widehat{\rho}\widehat{\tau}, \quad (7.214)$$

$$D(\widehat{\delta\mathcal{A}^2}) = \frac{1}{A} \left(2\widehat{\mathcal{A}^2}\widehat{\delta\rho} - 3\widehat{\rho}\widehat{\delta\mathcal{A}^2} + \overline{\widehat{\sigma}}\widehat{\delta\mathcal{A}^2} \right) - 2\widehat{\mathcal{A}^2}\widehat{\rho}\overline{\widehat{\tau}}, \quad (7.215)$$

where the divergent quantities are to be replaced by their regularised counterparts in $\delta\Psi_0$, $\overline{\delta\Psi_0}$, $\delta\Psi_1$ and $\overline{\delta\Psi_1}$, which are defined in (7.131)–(7.134).

Finally, we see from (7.202) that the quantity $\text{Re}(\mathbf{u}_1)$ scales like $(s-s_0)^{-2}$. We therefore introduced the regularised quantity $\text{Re}(\widehat{\mathbf{u}}_1) = A^2 \text{Re}(\mathbf{u}_1)$, which is not divergent at the caustic point. This quantity satisfies the transport equation

$$D(\text{Re}(\widehat{\mathbf{u}}_1)) = \frac{1}{A} \left[-4\widehat{\rho}\text{Re}(\widehat{\mathbf{u}}_1) + \frac{i}{2} \left(\widehat{\delta\chi} - (\widehat{\delta\chi})^* + \widehat{\chi}\frac{\widehat{\delta\mathcal{A}^2}}{\mathcal{A}^2} - \overline{\widehat{\chi}}\frac{\widehat{\delta\mathcal{A}^2}}{\mathcal{A}^2} \right) \right] \\ + \frac{i}{2} \left[\widehat{\rho}(\overline{\widehat{\mu}} - \widehat{\mu}) + \overline{\widehat{\sigma}}\overline{\widehat{\lambda}} - \widehat{\sigma}\widehat{\lambda} \right]. \quad (7.216)$$

We note that the regularisation operation (denoted by a caret) commutes with complex conjugation (denoted by an overline or an asterisk) because the cross-sectional area A is real. (Recall that a divergent quantity z which scales like $(s-s_0)^n$ at the caustic point is regularised by taking $\widehat{z} = A^{-n}z$. Then $\overline{\widehat{z}} = A^{-n}\overline{z} = \widehat{\overline{z}}$.)

7.8.2 Numerical method and results

Our principal aim is to consider spin–helicity effects for circularly polarised electromagnetic waves on Kerr spacetime. In particular, we shall consider a circularly polarised electromagnetic wave which begins far from the system, interacts with the black hole, and is then scattered to infinity. We employ the higher-order geometric optics formalism to consider the lensing effects of an initially circular infinitesimal bundle of rays in the electromagnetic wavefront. We shall evolve the transport equations of Section 7.8.1 along null geodesics of Kerr spacetime, in order to calculate $\text{Re}(\mathbf{u}_1)$ along a ray. (We consider the case of a rapidly rotating Kerr black hole of mass $M = 1$ with spin parameter $a = 0.9$.)

Null geodesic equations and initial conditions

We use the Hamiltonian formalism for null geodesics in Boyer–Lindquist coordinates. In particular, we evolve the full set of Hamilton’s equations, which is a system of eight first-order ordinary differential equations for the spacetime coordinates and their conjugate momenta. We recall that there are four conserved quantities along null rays in the Kerr spacetime: the energy $E = -p_t$, the azimuthal angular momentum $L_z = p_\phi$; the Hamiltonian $H = 0$; and the Carter constant $K = p_\theta^2 + (aE \sin \theta - L_z \operatorname{cosec} \theta)^2$.

Following Johannsen and Psaltis [9], we consider a ray which begins on an “image plane” whose centre is a distance D_0 away from the origin of the coordinate system centred on the black hole, at an inclination angle ι , with its centre in the (x, z) -plane; see Appendix of [9]. Since the ray begins far from the black hole, the initial data can be expressed using standard Euclidean geometry. The image plane coordinates (X, Y) are related to the Cartesian coordinates (x, y, z) centred on the black hole via

$$x = -Y \cos \iota + D_0 \sin \iota, \quad y = X, \quad z = Y \sin \iota + D_0 \cos \iota. \quad (7.217)$$

A photon which begins on the image plane at (X, Y) with initial three-momentum orthogonal to the image plane has initial conditions

$$r = \sqrt{X^2 + Y^2 + D_0^2}, \quad p_r = -\frac{D_0}{r}, \quad (7.218)$$

$$\theta = \arccos \left(\frac{Y \sin \iota + D_0 \cos \iota}{r} \right), \quad p_\theta = \frac{\cos \iota - D_0 (Y \sin \iota + D_0 \cos \iota)}{r^2 \sqrt{X^2 + (D_0 \sin \iota - Y \cos \iota)^2}}, \quad (7.219)$$

$$\phi = \arctan \left(\frac{X}{D_0 \sin \iota - Y \cos \iota} \right), \quad p_\phi = \frac{X \sin \iota}{X^2 + (D_0 \sin \iota - Y \cos \iota)^2}. \quad (7.220)$$

The t -component of the four-momentum is calculated from the three-momentum so that the norm $g^{ab} p_a p_b$ vanishes. In practice, we use the fact that $g_{ab} \rightarrow \eta_{ab}$ as $r \rightarrow \infty$ to calculate p_t . In this limit, the choice of normalisation of the three-momentum means that we have $p_t = -1$.

Transport equations and far-field initial data

In order to calculate $\operatorname{Re}(\mathbf{u}_1)$, one must evolve the system of transport equations (7.5) for the cross-sectional area; (7.3) for the square-amplitude; (7.100)–(7.102) for the Newman–Penrose scalars; (7.104)–(7.111) for the higher-order Newman–Penrose and geometrical optics quantities; and (7.12) for $\operatorname{Re}(\mathbf{u}_1)$ itself. These must be evolved along a central null geodesic, whose initial conditions are determined using the set-up presented in the previous section.

We wish to consider an (initially circular) infinitesimal bundle of rays which starts at infinity; however, in practice, the numerical integration of the transport equations must begin at some finite (but large) value of r . One may use the far-field solutions to the transport equations given in Section 7.6 to determine the leading-order-in- $\frac{1}{r}$ corrections to the initial conditions which arise from the fact that $r(0) = r_0 < \infty$.

For example, at infinity, one would choose the expansion ρ and shear σ to be zero initially. One can see from the Sachs equations (7.100) that, even in regions of spacetime

where the curvature is small, shear will be generated through the transport equation for σ , which will generate expansion through the transport equation for ρ . The leading-order correction to the initial conditions for ρ and σ are determined by (7.147), i.e.,

$$\rho(0) = \frac{1}{7E} \left(\frac{3KM}{4E} \right)^2 \frac{1}{r_0^7}, \quad \sigma(0) = \frac{3KM}{4E} \frac{1}{r_0^4}, \quad (7.221)$$

where $r_0 \gg 1$. The initial conditions for the other Newman–Penrose quantities can be determined using the leading-order far-field power series solutions given in (7.147)–(7.149).

We test the numerical method by integrating the transport equations along a ray with initial conditions

$$X(0) = 5, \quad Y(0) = 10, \quad D_0(0) = 10^5, \quad \iota(0) = \frac{17\pi}{18}. \quad (7.222)$$

The initial values of the spacetime coordinates and their conjugate momenta can be determined through (7.218) and (7.219). In particular, $r(0) = \sqrt{X(0)^2 + Y(0)^2 + D_0(0)^2}$. This value of $r(0)$ is then used to determine the initial data for the remaining quantities, using the far-field solutions of Section 7.6. We also choose $A(0) = 1$.

The transport equations can be evolved along the ray until we reach (the neighbourhood of) the first caustic point at $s = s_0$, where the Newman–Penrose quantities diverge; see Section 7.7. In practice, we can never reach the caustic point at $s = s_0$, so we must impose a halting condition: the numerical integration is stopped at $s = s_{-\varepsilon} = s_0 - \varepsilon$ (where $\varepsilon > 0$) when

$$A(s_{-\varepsilon}) = A_{-\varepsilon}, \quad 0 < A_{-\varepsilon} \ll 1; \quad (7.223)$$

typically, we take $A_{-\varepsilon} = 10^{-5}$. (Recall that $A(s_0) = 0$ at the caustic point, where neighbouring rays cross, and one of the two shape parameters d_{\pm} vanishes.)

In order to calibrate the numerical method and check the solutions to the transport equations, we fit a generalised power series solution to each of the variables in the vicinity of the caustic point, and compare with the results of Section 7.7. Here, we briefly outline the steps involved.

The output of the numerical integration is an interpolating function for each of the (complex) scalar quantities described above (i.e., the Newman–Penrose scalars, higher-order Newman–Penrose quantities, and geometric optics quantities). We take as numerical data the interpolating function evaluated at a selection of N equally spaced values of $s \in [s_{-\varepsilon} - s_*, s_{-\varepsilon}]$ within a distance s_* of the stopping value $s_{-\varepsilon}$. Here, s_* is taken to be some sufficiently small value, so that the generalised power series solution through sub-leading order is a good approximation for the true solution. (Typically, we take $s_* = 0.1$ and $N = 10^3$.) For a quantity which diverges like $(s - s_0)^n$ at leading order, we find a least-squares fit to the numerical data as a linear combination of the leading- and sub-leading-order functions, $(s - s_0)^n$ and $(s - s_0)^{n+1}$. (In practice, this is achieved using Mathematica’s `Fit` function.)

Upon inspection of plots in the neighbourhood of the caustic point, we see good agreement between the solutions to the transport equations and the least-squares fit through sub-leading order for all of the quantities under consideration.

Recall that, there are a number of free parameters in the sub-leading-order generalised power series solutions obtained in Section 7.7. The free parameters which arise from the leading- and sub-leading-order solutions to the transport equations for the Newman–Penrose scalars $\{\rho, \sigma, \chi, \tau, \mu, \lambda\}$ are $\varphi \in [0, 2\pi)$, $\rho_1, |\chi_0|, |\chi_1|, |\tau_0|, |\tau_1| \in \mathbb{R}$ and $\mu_0, \mu_1 \in \mathbb{C}$. The cross-sectional area A features one free parameter $A_1 \in \mathbb{R}$ which appears at leading order. The square-amplitude \mathcal{A}^2 also features one free parameter $a_0 \in \mathbb{R}$ which arises from the leading-order piece. The free parameters which appear in the higher-order Newman–Penrose quantities $\{\delta\rho, \bar{\delta}\rho, \delta\sigma, \bar{\delta}\sigma, \delta\chi, \bar{\delta}\chi\}$ are $\alpha_0 \in \mathbb{R}$ which comes from the leading-order part, and $\xi_1, \eta_1 \in \mathbb{R}$ which arise from the sub-leading-order part. Finally, in the higher-order geometric optics quantity $\delta\mathcal{A}^2$, there is one free parameter $\zeta_1 \in \mathbb{R}$ which arises at sub-leading order. For each of the quantities described above, we may extract the free parameters which appear in the coefficients of the generalised power series solutions using the least-squares fits through sub-leading order. We see consistency between the free parameters which are extracted from the numerical solutions.

Evolving regularised transport equations through caustic points

In order to calculate quantities along a ray beyond the first caustic point, we use the regularised transport equations of Section 7.8.1. Here, we briefly outline the numerical method, and illustrate this by evolving the regularised transport equations along a reference trajectory.

The full system of equations we evolve is given by the geodesic equations (i.e., Hamilton’s equations) for the four spacetime coordinates and their conjugate momenta, as well as the set of regularised transport equations (7.203)–(7.216). The right-hand sides of the regularised equations feature denominators which are singular at the caustic point. It is therefore necessary to integrate just short of the caustic point, and take a “step” over the caustic point, before resuming the numerical integration on the other side.

The regularised quantities are all related to the original quantities by a factor of A to some power. Since we choose to set $A(0) = 1$, the initial data for the regularised quantities is the same as that for the original (divergent) quantities described above. We integrate the full system of ordinary differential equations along a central ray with initial data given by (7.222) until we reach the neighbourhood of the first caustic point. In practice, we integrate from $s = 0$ up to $s = s_{-\varepsilon} = s_0 - \varepsilon$, where $A(s_{-\varepsilon}) = A_{-\varepsilon}$ for some predetermined tolerance parameter $0 < A_{-\varepsilon} \ll 1$; see the halting condition (7.223). Performing a Taylor expansion of $A(s)$ in the neighbourhood of the stopping point $s = s_{-\varepsilon}$, we find

$$A(s) \sim A_{-\varepsilon} - 2\widehat{\rho}_{-\varepsilon}(s - s_{-\varepsilon}), \quad (7.224)$$

where $\widehat{\rho}_{-\varepsilon} = \widehat{\rho}(s_{-\varepsilon})$, and we have used the regularised transport equation (7.206) to replace the coefficient at $O(s - s_{-\varepsilon})$. Evaluating (7.224) at the caustic point $s = s_0$, noting that $A(s_0) = 0$ by definition, and rearranging, we find

$$\varepsilon = \frac{A_{-\varepsilon}}{\widehat{\rho}_{-\varepsilon}}. \quad (7.225)$$

Hence, given the location of the stopping point $s_{-\varepsilon}$, one can approximate the location of

the caustic point $s_0 = s_\varepsilon + \varepsilon$ using information about $A(s)$ and its first derivative – i.e., $\widehat{\rho}(s)$ – at the stopping point.

To evolve the transport equations “through” the caustic point, we take a step of size 2ε from the stopping point to $s_\varepsilon = s_0 + \varepsilon$ and restart the numerical integration of the regularised transport equations and Hamilton’s equations. Beyond the caustic point, at $s = s_\varepsilon$, the new initial data are determined using information about the first derivative of each quantity at the stopping point. As an illustrative example, consider the evolution of $\widehat{\rho}(s)$, which satisfies the transport equation $D\widehat{\rho} = A^{-1}(\widehat{\sigma}\overline{\widehat{\sigma}} - \widehat{\rho}^2)$. The value of $\widehat{\rho}$ at the stopping point is $\widehat{\rho}_{-\varepsilon} = \widehat{\rho}(s_{-\varepsilon})$. Using the Taylor series expansion of $\widehat{\rho}(s)$ about the stopping point, one may write

$$\widehat{\rho}(s) \sim \widehat{\rho}_{-\varepsilon} + D\widehat{\rho}|_{s=s_{-\varepsilon}}(s - s_{-\varepsilon}), \quad (7.226)$$

where it is understood that the term $D\widehat{\rho}$ is to be replaced by the right-hand side of the transport equation for $\widehat{\rho}$. Evaluating (7.226) at $s = s_\varepsilon$, the (first-order approximation to the) new initial condition after the caustic point is

$$\widehat{\rho}(s_\varepsilon) = \widehat{\rho}_{-\varepsilon} + \frac{2\varepsilon}{A_{-\varepsilon}} \left(\widehat{\sigma}_{-\varepsilon}\overline{\widehat{\sigma}_{-\varepsilon}} - \widehat{\rho}_{-\varepsilon}^2 \right), \quad (7.227)$$

where $\widehat{\sigma}_{-\varepsilon} = \widehat{\sigma}(s_{-\varepsilon})$ and ε is determined by (7.225). The initial data for all other quantities of interest may be determined in a similar fashion. Once the new initial data are known, the ordinary differential equations may be evolved for $s \geq s_\varepsilon$. In practice, we carry out the numerical integration until either $\Delta(r) = r^2 - 2Mr + a^2 \leq \varepsilon_H \ll 1$, or $r \geq R \gg 1$. The former case corresponds to a ray which approaches the outer event horizon; the latter case corresponds to a ray which reaches infinity. (The particular ray we consider here satisfies the second condition.)

Once the regularised quantities have been obtained by numerically solving the transport equations (7.203)–(7.216), the divergent quantities can be obtained straightforwardly, e.g. $\rho = A^{-1}\widehat{\rho}$. In order to check the validity of the numerical solutions to the transport equations, one may compare them to the generalised power series solutions through sub-leading order. The latter are valid on *both* sides of the caustic point (for $|s - s_0|$ sufficiently small), so a comparison between the two solutions permits us to check whether the evolution through the caustic point has been successful.

In Figures 7.1–7.3, we present the numerical results obtained by evolving the regularised transport equations alongside the sub-leading-order generalised power series solutions in the vicinity of the caustic point. For the particular ray we consider, we set $A_{-\varepsilon} = 10^{-5}$. The value of the affine parameter when the halting condition is imposed is $s_{-\varepsilon} = 100023.582612$, and the (estimated) value of the affine parameter at the caustic is $s_0 = 100023.582715$. Let us now describe the results of Figures 7.1–7.3 in more detail.

Figure 7.1 shows the Newman–Penrose scalars $\{\rho = A^{-1}\widehat{\rho}, \sigma = A^{-1}\widehat{\sigma}, \chi = A^{-1}\widehat{\chi}, \tau = A^{-1}\widehat{\tau}, \mu = A^{-1}\widehat{\mu}, \lambda = A^{-1}\widehat{\lambda}\}$, which are built from the regularised quantities, in the neighbourhood of the caustic point at $s = s_0$. The real (imaginary) part of the numerical solution to the transport equations is plotted in blue (orange). As described above, the quantities are evolved from $s = 0$ until $s = s_{-\varepsilon}$, where $A(s) = A_{-\varepsilon} \ll 1$. The location of the caustic point, $s = s_0$, is estimated by extrapolation of $A(s)$. The numerical integration

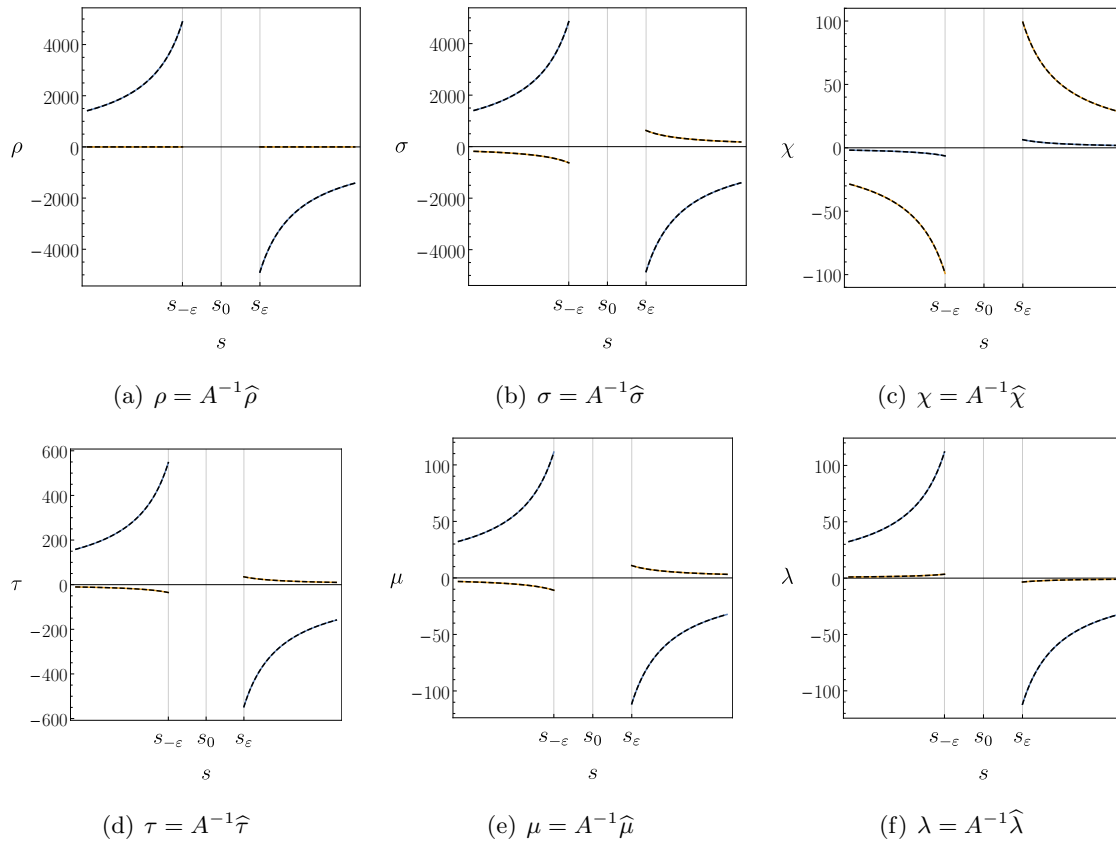


Figure 7.1. Newman–Penrose scalars in the neighbourhood of the caustic point of multiplicity one at $s = s_0$, where s denotes the affine parameter. The real (imaginary) part of each Newman–Penrose scalar is shown in blue (orange). The generalised power series solutions through sub-leading order are overlaid as black dashed curves. We see good agreement between the interpolating functions and the generalised power series expansions through sub-leading order on both sides of the caustic point. (The absolute relative error is always less than 10^{-4} .)

is then restarted at $s = s_\varepsilon = s_{-\varepsilon} + 2\varepsilon$ using the initial conditions described above. The generalised power series solutions (7.164)–(7.166) are shown as black dashed curves. We see good agreement between the numerical data and the least-squares fits of the generalised power series through sub-leading order on both sides of the caustic point: the absolute value of the relative error between the numerical solution and the least-squares fit is always less than 10^{-4} . We note that the solutions are odd about $s = s_0$ because the leading-order terms at $O((s - s_0)^{-1})$ are equal and opposite on either side of $s = s_0$.

In Figure 7.2 we present numerical results for the higher-order Newman–Penrose quantities $\{\delta\rho = A^{-3}\widehat{\delta\rho}, \bar{\delta}\rho = A^{-3}\widehat{\bar{\delta}\rho}, \delta\sigma = A^{-3}\widehat{\delta\sigma}, \bar{\delta}\sigma = A^{-3}\widehat{\bar{\delta}\sigma}, \delta\chi = A^{-3}\widehat{\delta\chi}, \bar{\delta}\chi = A^{-3}\widehat{\bar{\delta}\chi}\}$ in the neighbourhood of the caustic point at $s = s_0$. Again, the real (imaginary) parts of the solutions to the regularised transport equations are shown in blue (orange). The least-squares fits of the sub-leading-order generalised power series solutions (7.190)–(7.195) are shown as black dashed curves. We again see good agreement with the numerical solutions to the transport equations; the absolute value of the relative error is always less than 10^{-4} . We clearly see that there is an asymmetry in the solutions about $s = s_0$, due to

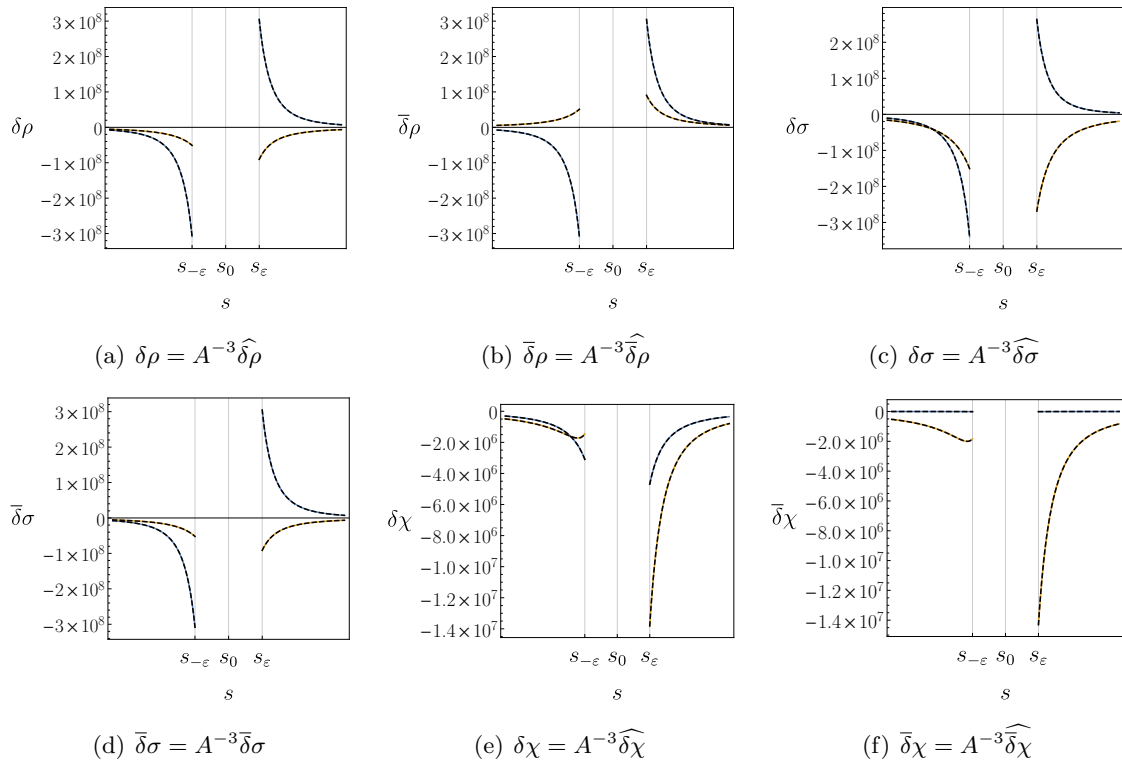


Figure 7.2. Higher-order Newman–Penrose quantities in the neighbourhood of the caustic point of multiplicity one at $s = s_0$. The real (imaginary) part of each quantity is shown in blue (orange). The generalised power series solutions through sub-leading order are overlaid as black dashed curves. There is good agreement between the interpolating functions and the generalised power series expansions through sub-leading order on both sides of the caustic point; the absolute relative error is always less than 10^{-4} .

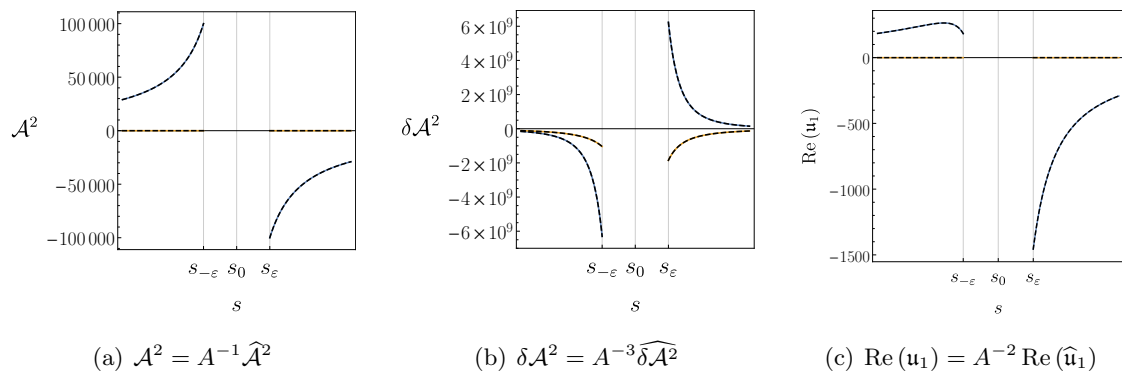


Figure 7.3. Geometric optics quantities in the neighbourhood of the caustic point of multiplicity one at $s = s_0$. The real (imaginary) part of each quantity is shown in blue (orange). The generalised power series solutions through sub-leading order are overlaid as black dashed curves. We observe good agreement between the interpolating functions and the generalised power series expansions through sub-leading order on both sides of the caustic point.

the presence of competing terms at leading and sub-leading order, i.e., $O((s - s_0)^{-3})$ and $O((s - s_0)^{-2})$. Intriguingly, we find that, for certain quantities, the sub-leading-order terms dominate until very close to the caustic point, due to the fact that the leading-order coefficients are small in magnitude.

Figure 7.3 shows the geometric optics quantities $\mathcal{A}^2 = A^{-1}\widehat{\mathcal{A}}^2$, $\delta\mathcal{A}^2 = A^{-3}\widehat{\delta\mathcal{A}}^2$ and $\text{Re}(\mathbf{u}_1) = A^{-2}\text{Re}(\widehat{\mathbf{u}}_1)$. We again see good agreement between the interpolating functions built from numerical solutions to the regularised transport equations and the sub-leading order generalised power series solutions (7.169) and (7.200). For each of these quantities, the absolute value of the relative error is less than 10^{-4} in the neighbourhood of the caustic point.

Recall that the regularisation method involves halting the numerical integrator just short of the caustic point at $s = s_{-\varepsilon}$, where $A(s_{-\varepsilon}) = A_{-\varepsilon}$. We estimate the value of the affine parameter at the caustic point ($s = s_0$) via extrapolation; see (7.224) and (7.225). Dolan [81, 171] demonstrated that certain Newman–Penrose quantities (e.g. ρ and σ) may be calculated beyond caustic points using a second-order formalism; see Section 3.3.2 for a review. The complex scalar quantity c , which corresponds to a point in the instantaneous geometric optics wavefront, satisfies the second-order equation (3.43) which is well-behaved at caustic points. We recall that the shape parameters of the bundle satisfy $d_+d_- = 2|\text{Im}(c_1\bar{c}_2)|$, where c_1 and c_2 are any pair of linearly independent solutions to (3.43). At a caustic point, one of the shape parameters d_{\pm} vanishes. One may therefore calculate the value of the affine parameter at the caustic point independently by solving the equation $\text{Im}(c_1\bar{c}_2) = 0$ for s . Doing this for our reference ray, we find that the absolute value of the relative error between the value estimated via extrapolation of the regularised transport equations and the value obtained using the second-order equation (3.43) is approximately 10^{-11} .

Sensitivity to internal parameters

Let us now quantify the sensitivity of the numerical results to changes in the numerical method’s internal parameters, such as the caustic-handling parameter $A_{-\varepsilon}$. To do this, we twice solve the full system of transport equations subject to the initial data described above, comparing the results obtained for $A_{-\varepsilon}^{(1)} = 10^{-5}$ and $A_{-\varepsilon}^{(2)} = 2 \times 10^{-5}$.

We find that the Newman–Penrose scalars are insensitive to changes in the caustic-handling parameter $A_{-\varepsilon}$ beyond the caustic point. The absolute value of the relative error between the two runs is less than approximately 10^{-5} for all regularised Newman–Penrose scalars (e.g. $\widehat{\rho}$ and $\widehat{\sigma}$) beyond the caustic point. We can therefore be confident that the regularisation method and its numerical implementation have allowed us to calculate the Newman–Penrose scalars beyond the caustic point, far from the black hole.

In addition, we observe good agreement for the higher-order quantities, such as $\widehat{\delta\rho}$, $\widehat{\delta\chi}$ and $\text{Re}(\widehat{\mathbf{u}}_1)$, in the neighbourhood of the caustic point: within a distance $s' = s - s_0 \sim 0.005$ of the caustic point, the absolute value of the relative error between the two runs is $O(10^{-4})$ or smaller. However, we find that the higher-order quantities become highly sensitive to changes in the caustic-handling parameter downstream from the caustic point (i.e., as s' is increased). A consequence is that we cannot be confident in the numerical results for

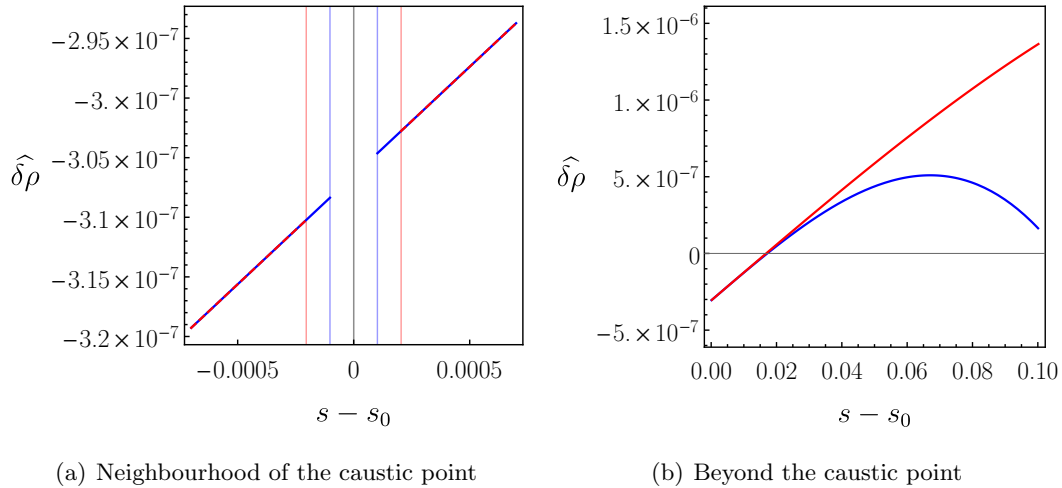


Figure 7.4. Numerical results for $\text{Re}(\widehat{\delta\rho})$ for two runs using $A_{-\varepsilon}^{(1)} = 10^{-5}$ [blue] and $A_{-\varepsilon}^{(2)} = 2 \times 10^{-5}$ [red]. The horizontal axis is rescaled, so that the caustic point occurs at $s' = s - s_0 = 0$. (a) Numerical solutions for the two runs in the vicinity of the caustic point. We observe good agreement between the curves obtained from the first [blue] and second [red, dashed] runs. (b) Disagreement between the two solutions beyond the caustic point. The curves begin to diverge for $s' = s - s_0 \sim 0.02$.

the higher-order quantities far from the black hole. For example, when $A_{-\varepsilon} = A_{-\varepsilon}^{(1)}$, we find $\lim_{s \rightarrow \infty} \text{Re}(\mathbf{u}_1) \sim 200$; however, when $A_{-\varepsilon} = A_{-\varepsilon}^{(2)}$, we obtain $\lim_{s \rightarrow \infty} \text{Re}(\mathbf{u}_1) \sim -30$.

Figure 7.4 shows the numerical results for $\text{Re}(\widehat{\delta\rho})$ – chosen as an example of a typical (regularised) higher-order Newman–Penrose quantity – with $A_{-\varepsilon} = A_{-\varepsilon}^{(1)}$ [blue] and $A_{-\varepsilon} = A_{-\varepsilon}^{(2)}$ [red]. In both figures, the horizontal axis has been rescaled ($s \mapsto s' = s - s_0$) so that the caustic point occurs at $s' = 0$. In Figure 7.4(a), we compare the results in the vicinity of the caustic point. For the first run, the numerical integration is halted when $A = A_{-\varepsilon}^{(1)}$; the corresponding value of $s_{-\varepsilon}^{(1)}$ is shown as a blue vertical line. The value $s_0^{(1)} = s_{-\varepsilon}^{(1)} + \varepsilon^{(1)}$, which is estimated via extrapolation using (7.225), is shown as a vertical grey line. The numerical integration is resumed at $s_{\varepsilon}^{(1)} = s_0^{(1)} + \varepsilon$, shown using a second vertical blue line. The second run is halted when $A = A_{-\varepsilon}^{(2)}$. The corresponding values of $s_{\pm\varepsilon}^{(2)}$ are shown as red vertical lines. The estimated value of the affine parameter at the caustic point, denoted $s_0^{(2)}$, is *not* identical to the value obtained from the first run; however, the absolute difference between these two quantities is $O(10^{-10})$. The numerical results for $\text{Re}(\widehat{\delta\rho})$ are in good agreement on both sides of the caustic point: on the domain shown in Figure 7.4(a), the absolute value of the relative error between the numerical data obtained from the two runs is less than $O(10^{-5})$. Figure 7.4(b) shows the numerical results for $\text{Re}(\widehat{\delta\rho})$ beyond the caustic point (which is again located at $s' = 0$). One can clearly see that for $s' \sim 0.02$, the curves begin to diverge from one another. We observe similar behaviour for all other higher-order quantities. This demonstrates that the numerical solutions to the transport equations beyond the caustic point are highly sensitive to small changes in the caustic-handling parameter $A_{-\varepsilon}$. In Section 7.9.1, we comment on potential ways that the issues described here could be alleviated.

7.9 Discussion

7.9.1 Extensions

Construction of parallel-transported complex null tetrads

In Section 7.4.1, we constructed a complex null tetrad by availing the properties of the principal tensor h_{ab} and its Hodge dual f_{ab} . The construction of the parallel-transported complex null tetrad should therefore apply to any geometry which admits a principal tensor, i.e., the (off-shell) Kerr–NUT–(anti-)de Sitter spacetime. The parallel-transported complex null tetrad constructed here can be of practical use in other scenarios beyond gravitational lensing by Kerr black holes.

We recall that the construction of the complex null tetrads of Section 7.4.1 breaks down for rays with $K = 0$. This is due to the fact that the tangent vector k^a and the vector $f^a_b k^b$ are collinear in this degenerate case. One could extend the work of Section 7.4.1 by paying special attention to the construction of a parallel-propagated complex null tetrad along null geodesics with vanishing Carter constant.

Modified geometric optics, transport equations and caustic points

The results of Figure 7.4 demonstrate that the numerical solutions to the transport equations of modified geometric optics are sensitive to small changes in the caustic-handling parameter $A_{-\varepsilon}$ beyond the caustic point. This hints at a potential flaw in either (i) the modified geometric optics formalism (Section 7.2); (ii) its implementation in terms of first-order transport equations (Section 7.5); or (iii) the numerical method used to evolve transport equations through caustic points (Section 7.8).

Let us first consider point (i). The geometric optics approximation relies on the assumption that the wavelength is small in comparison to other physically relevant length scales (e.g. spacetime curvature scales). This assumption breaks down at caustic points. The geometric optics approximation becomes valid again after passing through a caustic point, when the rays in the bundle are sufficiently separated. However, the passage through the caustic point gives rise to a change in sign of the Jacobian of the lensing map [318, 319]. The (complex) geometric optics amplitude will therefore change phase by $\pm\frac{\pi}{2}$ through a caustic, becoming purely imaginary ($\mathcal{A} \mapsto \pm i\mathcal{A}$). Zenginoğlu and Galley [319] remark that the sign of the phase shift cannot be computed within the geometric optics approximation; instead, this must be determined by matching to the full solution from numerical simulations, or by some alternative approach. It is clear that a better understanding of caustic points and the passage through them is required in order to overcome the issues with the modified geometric optics formalism.

Now consider point (ii) from the above list. The leading-order geometric optics formalism (Section 7.2.1) and its sub-leading-order modification (Section 7.2.2) are formulated in terms of first-order transport equations for quantities which diverge at caustic points. This issue could perhaps be alleviated by adopting an alternative approach. Recall that a subset of the Newman–Penrose scalars can be obtained beyond caustic points by solving a system of first- and second-order ordinary differential equations along rays; see Section

3.3.2 and [81, 171]. The quantities which satisfy these differential equations do not diverge, and can be evolved through caustic points without issue. It is not clear whether there exists a similar system of equations satisfied by well-behaved quantities from which higher-order Newman–Penrose quantities (e.g. $\delta\rho$) and geometric optics quantities (e.g. \mathbf{u}_1) may be obtained.

Finally, let us consider point (iii) from the above list. There are a number of ways that the accuracy of the numerical method presented in Section 7.8.2 could be improved. First, recall that we use extrapolation to step over the caustic point, based on a first-order Taylor expansion of the dependent variables; see (7.224) and (7.226), for example. For improved accuracy, one could extend this Taylor expansion to higher orders; this would involve taking further derivatives of the dependent variables in order to work out the new initial data after the caustic point. Alternatively, one could use a “midpoint method” to calculate the initial data at $s = s_\varepsilon$, taking an average of the values obtained by performing (first-order) Taylor series expansions about $s = s_{-\varepsilon}$ and $s = s_\varepsilon$. Second, recall that we obtained generalised power series solutions through sub-leading order for all of the Newman–Penrose scalars and higher-order Newman–Penrose quantities in the vicinity of a caustic point of multiplicity one (Section 7.7). These generalised power series solutions are valid on both sides of the caustic point; we therefore used them as a consistency check of the numerical solution obtained by regularisation beyond the caustic point; see Figures 7.1–7.3. In addition to employing these solutions as a consistency check, one could use them to perform a more sophisticated evolution of the divergent quantities through the caustic point, by performing “matched asymptotic expansions” in the neighbourhood of the caustic point.

In a recent paper, Dolan [320] presented a method to calculate the vector potential for electromagnetic waves on Kerr spacetime in Lorenz gauge. To further test the modified geometric optics formalism, one could perform a direct comparison between the results obtained from solving wave equations in Lorenz gauge, and those obtained by evolving systems of transport equations (considered here).

7.9.2 Conclusions

In this chapter, we have employed the higher-order geometric optics formalism of Dolan [81] (see Section 7.2) to understand aspects of gravitational lensing by Kerr black holes in general relativity. Our principal aim was to calculate the $O(\omega^{-1})$ correction to the electromagnetic stress–energy tensor (7.11), which involves the quantities \mathbf{u}_1 , \mathbf{v}_1 and \mathbf{w}_1 . In Section IV of [81], Dolan enumerates several practical hurdles which one faces when applying the higher-order geometric optics formalism, including: (1) finding a suitable parallel-transported complex null tetrad; (2) calculating key quantities, such as the Weyl curvature scalars and their derivatives; (3) solving transport equations numerically or otherwise; and (4) handling ray-crossings and caustic points.

For Kerr spacetime (reviewed in Section 7.3), we used the complex bivector (7.54) to construct a complex null tetrad $\{k^a, \tilde{n}^a, \tilde{m}^a, \overline{\tilde{m}}^a\}$ along null geodesics by projecting along the tangent vector k^a . Performing a Lorentz transformation (7.66) with complex rotation parameter B given by (7.67), we mapped the “tilded” tetrad onto a complex null tetrad

$\{k^a, n^a, m^a, \bar{m}^a\}$ which is parallel-transported along the null geodesics of Kerr spacetime. Thus we have overcome hurdle (1) from the above list. Intriguingly, we were able to arrive at this tetrad by defining a new operator (7.69) and projecting along k^a , bypassing the “tilded” tetrad altogether. Moreover, we demonstrated that the parallel-propagated tetrad constructed here is equivalent to a complex null tetrad constructed from the legs of Marck’s (real) parallel-transported tetrad [52].

In order to calculate the quantities \mathbf{u}_1 , \mathbf{v}_1 and \mathbf{w}_1 in the sub-leading-order correction to the electromagnetic stress–energy tensor, it is necessary to evolve a system of transport equations for Newman–Penrose quantities along rays (see Section 7.5.1). These transport equations feature the complex Weyl scalars Ψ_i and their derivatives along the wavefront (e.g. $\delta\Psi_0$). The Weyl scalars take a simple form in the “tilded” tetrad, and can be transformed to the parallel-transported tetrad straightforwardly by means of the Lorentz transformation (7.66); see Section 7.5.2. The calculation of the quantities $\{\delta\Psi_0, \bar{\delta}\Psi_0, \delta\Psi_1, \bar{\delta}\Psi_1\}$ is more involved; however, progress can be made by transforming to the symmetrised Kinnersley tetrad. Explicit expressions for these quantities are given for the parallel-propagated tetrad in Appendix E. Hence, we have successfully calculated closed-form expressions for the key quantities described in point (2) of the above list. In Section 7.6, we analysed the far-field behaviour of these quantities, and used the results to understand the asymptotic behaviour of the Newman–Penrose scalars for $r \rightarrow \infty$, which is useful in gravitational lensing calculations.

Points (3) and (4) from the list pose more of an issue. In principle, the relevant transport equations (Section 7.5.1) can be calculated numerically along a ray. However, on Kerr spacetime (and indeed in other situations), caustic points arise generically. The geometric optics formalism breaks down at caustic points, and quantities like ρ (expansion) and σ (shear) diverge there. In Sections 7.7.2 and 7.7.3, we used the full system of Newman–Penrose transport equations (Section 7.5.1) to deduce the behaviour of Newman–Penrose scalars and higher-order geometric optics quantities in the neighbourhood of a caustic point of multiplicity one. In particular, we obtained these quantities as generalised power series in $(s - s_0)$ through sub-leading order, where $s = s_0$ is the value of the affine parameter at the caustic point, as measured along the central ray of the bundle. The results of Section 7.7.2 are general, in that they are valid for any Ricci-flat manifold.

In Section 7.8, we presented a practical method to evolve the full system of Newman–Penrose transport equations through caustic points. The method involves regularising divergent quantities at caustic points: a complex scalar quantity z which scales like $(s - s_0)^n$ at a caustic point is regularised by taking $\hat{z} = A^{-n}z$. The regularised variables are finite at caustic points. However, the transport equations feature denominators which are singular at the caustic point. We therefore integrate the transport equations just short of the caustic point, then use a first-order Taylor expansion of the regularised quantities to “step over” the caustic point. The integration is then resumed on the other side of the caustic point. We presented some preliminary numerical results for a reference trajectory on Kerr spacetime (with $M = 1$, $a = 0.9$). In Figures 7.1–7.3, we observe good agreement between the numerical solution to the transport equations and a least-squares fit of the sub-leading-order generalised power series solutions in the neighbourhood of the caustic

point. Moreover, we find that the regularised Newman–Penrose scalars (e.g. $\hat{\rho}$ and $\hat{\sigma}$) are insensitive to changes in caustic-handling parameter $A_{-\varepsilon}$ downstream of the caustic point. This suggests that the results obtained for the Newman–Penrose scalars are valid far from the black hole, after scattering has occurred. However, we find that the higher-order quantities (e.g. $\hat{\delta}\rho$ and $\text{Re}(\hat{\mathbf{u}}_1)$) are sensitive to small changes in $A_{-\varepsilon}$ beyond the caustic point, as shown in Figure 7.4. We comment on this issue in Section 7.9.1.

Appendices

Appendix A

Translation between symbolic codes

Recall from Chapter 4 that the symbolic dynamics of Cornish and Gibbons [213], which we refer to as *collision dynamics*, may be used to describe the perpetual orbits of the Majumdar–Papapetrou di-hole in the meridian plane. This is achieved using the decision dynamics alphabet $\mathcal{A}_{CG} = \{+1, 0, -1\}$. In Section 4.5.3, we present an alternative symbolic code – *decision dynamics* – to describe trajectories in the Majumdar–Papapetrou di-hole geometry. This is achieved using the symbolic alphabet $\mathcal{A} = \{0, 1, 2, 3, 4\}$. The perpetual orbits can be described using infinitely long sequences from the reduced alphabet $\mathcal{A}_R = \{0, 2, 4\}$.

Consider a perpetual orbit described by the sequence $a = a_1 a_2 a_3 \cdots$, with $a_i \in \mathcal{A}_{CG}$, $\forall i \in \mathbb{N}$, as described by Cornish and Gibbons’s collision dynamics. Given the sequence a in collision dynamics, we would like to be able to translate to a sequence $\mathcal{X} = \mathcal{X}_1 \mathcal{X}_2 \mathcal{X}_3 \cdots$, with $\mathcal{X}_i \in \mathcal{A}_R$, $\forall i \in \mathbb{N}$, as described by decision dynamics. This can be achieved using the following algorithm.

We note first that, in decision dynamics, a “decision point” is preceded by the symbol ± 1 in collision dynamics. Such points are referred to as *pivots*, and are central to the translation between the two symbolic codes. We assume that $a_1 = \pm 1$ is our first pivot. Consider the first decision point, which follows the pivot $a_1 = \pm 1$. If $|a_1 - a_2| = 2$ (i.e., a_1 and a_2 have opposite signs), then replace a_1 with $\mathcal{X}_1 = 0 \in \mathcal{A}_R$ in decision dynamics. We then use $a_2 = \mp 1$ as the new pivot. On the other hand, if $|a_1 - a_2| = 1$ (i.e., $a_2 = 0$), we must consider the third element a_3 of the collision dynamics sequence. If $|a_1 - a_3| = 2$ (i.e., a_1 and a_3 have opposite signs), then we replace $a_1 a_2$ with $\mathcal{X}_1 = 2 \in \mathcal{A}_R$ in decision dynamics. Alternatively, if $|a_1 - a_3| = 0$ (i.e., a_1 and a_3 have the same sign), then we replace $a_1 a_2$ with $\mathcal{X}_1 = 4 \in \mathcal{A}_R$ in decision dynamics. In both of these cases, we use $a_3 = \pm 1$ as the new pivot. The whole sequence can be translated from collision dynamics to decision dynamics by continuing in this fashion, all the way along the sequence from left to right.

The algorithm is summarised in Table A.1. The leftmost column, denoted a_n , represents the pivot under consideration; this is always $a_n = \pm 1$ in collision dynamics. The second and third columns correspond to the subsequent elements, denoted by a_{n+1} and

a_n	a_{n+1}	a_{n+2}	\mathcal{X}_n
± 1			
	∓ 1		0
	0	∓ 1	2
	0	± 1	4

Table A.1. Summary of the translation between collision dynamics and decision dynamics. Beginning with the pivot $a_n = \pm 1$, we replace a finite substring in collision dynamics – either a_n or $a_n a_{n+1}$ – with a symbol \mathcal{X}_n from decision dynamics. If a_n and a_{n+1} have opposite signs, then we make the replacement $a_n \mapsto \mathcal{X}_n = 0$, and we use $a_{n+1} = \mp 1$ as the new pivot. If $a_{n+1} = 0$, then we consider a_{n+2} . In this case, if a_n and a_{n+2} have the opposite (same) sign, then we make the replacement $a_n a_{n+1} \mapsto \mathcal{X}_n = 2$ ($a_n a_{n+1} \mapsto \mathcal{X}_n = 4$), and we use $a_{n+2} = \pm 1$ as the new pivot.

a_{n+2} , of the finite substring we wish to translate. The final column, labelled \mathcal{X}_n , gives the corresponding symbol from the reduced decision dynamics alphabet \mathcal{A}_R . In each case, the final element of the finite substring under consideration – be that a_{n+1} or a_{n+2} – becomes the new pivot; this always takes the value ± 1 .

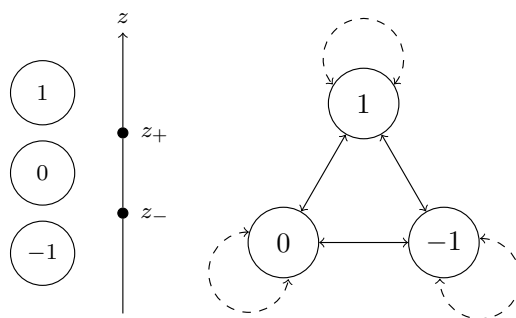


Figure A.1. Collision dynamics for the Majumdar–Papapetrou di-hole. The left-hand figure shows a schematic diagram of the symbolic coding. The right-hand figure shows the transition diagram for this symbolic dynamics: allowed transitions are shown using a solid line; forbidden transitions are shown using a dashed line. (Adapted from Figure 5 of [213].)

Figure A.1 shows a schematic diagram of the collision dynamics symbolic coding of Cornish and Gibbons [213]. The two black holes are located along the z -axis at $z = z_{\pm}$. Each time a null geodesic passes through $z \in (z_+, \infty)$ we assign the digit 1; when the ray passes through $z \in (z_-, z_+)$ we assign the symbol 0; and when the ray passes through $z \in (-\infty, z_-)$ we assign the digit -1 . The transition diagram on the right-hand side of Figure A.1 encodes the fact that no digit can follow itself. This highlights explicitly the correspondence with the Gaspard–Rice three disc model (see Section 4.5.2).

Appendix B

Circular photon orbits of the Majumdar–Papapetrou di-hole

In this section, we derive some of the key results of Chapter 4 and Chapter 6 on the existence of stable photon orbits around the equal-mass Majumdar–Papapetrou di-hole. These results were originally presented in [201].

In this analysis, we use units in which the individual black hole masses are set to unity, i.e., $M_{\pm} = M = 1$. This is equivalent to a rescaling of the coordinates $q^a \mapsto \frac{1}{M}q^a$, momenta $p_a \mapsto \frac{1}{M}p_a$, and the separation $d \mapsto \frac{d}{M}$.

B.1 Equatorial circular photon orbits

As described in Sections 4.3.2 and 6.3.1, null geodesic motion for the equal-mass Majumdar–Papapetrou di-hole is governed by the two-dimensional effective potential

$$h(\rho, z) = \rho U^2, \quad U(\rho, z) = 1 + \frac{1}{\sqrt{\rho^2 + (z - \frac{d}{2})^2}} + \frac{1}{\sqrt{\rho^2 + (z + \frac{d}{2})^2}}. \quad (\text{B.1})$$

The equal-mass di-hole is symmetric about the equatorial ($z = 0$) plane. Motion confined to this surface is governed by a one-dimensional effective potential $\hat{h}(\rho) = h(\rho, 0) = \rho \hat{U}^2$, where $\hat{U}(\rho) = 1 + \frac{2}{R}$, where we have introduced a new radial coordinate $R(\rho) = \sqrt{\rho^2 + \frac{d^2}{4}}$. (Note that $R > \frac{d^2}{4} > 0$ and $\hat{U} > 0$.) Circular null geodesics exist where $\hat{h}_{,\rho} = 0$. Differentiation of the effective potential gives

$$\hat{h}_{,\rho} = \frac{\hat{U}}{R^3} P(R), \quad P(R) = R^3 - 2R^2 + d^2. \quad (\text{B.2})$$

The positivity of \hat{U} and R mean that the task of locating circular photon orbits reduces to finding the roots of the cubic $P(R) = R^3 - 2R^2 + d^2$: circular orbits exist where $P(R)$ has roots. The discriminant of this polynomial is

$$\Delta_R(P) = -d^2 (27d^2 - 32). \quad (\text{B.3})$$

For $d > d_2 = \sqrt{\frac{32}{27}}$, $\Delta_R(P) < 0$, so $P(R)$ has one real root. It is straightforward to show that, in this regime, there are no roots with $R > 0$, and hence no equatorial circular

photon orbits. If $d = d_2$, $\Delta_R(P) = 0$, and there is a repeated root with $R > 0$. In this case there is one circular photon orbit. For $0 < d < d_2$, we have $\Delta_R(P) > 0$, so $P(R)$ has three distinct real roots, two of which have $R > 0$. Thus, there are two circular photon orbits in the equatorial plane.

The stability of these orbits can be determined by considering $\hat{h}_{,\rho\rho}$: orbits are stable (unstable) under radial perturbations if $\hat{h}_{,\rho\rho} < 0$ ($\hat{h}_{,\rho\rho} > 0$). In the case where $d < d_2$, the inner (outer) orbit is stable (unstable) under radial perturbations. In the critical case $d = d_2$, the unique circular photon orbit is marginally stable ($\hat{h}_{,\rho\rho} = 0$).

One may also consider the stability of equatorial circular photon orbits under perturbations in the z -direction. The equatorial symmetry of the equal-mass Majumdar–Papapetrou di-hole implies that $U_{,z}(\rho, 0) = 0$, thus $h_{,z}(\rho, 0) = 0$. Taking a second derivative, we find

$$h_{,zz}(\rho, 0) = 2\rho\hat{U}U_{,zz}(\rho, 0), \quad U_{,zz}(\rho, 0) = \frac{3d^2 - 4R^2}{2R^5}. \quad (\text{B.4})$$

One can deduce the stability of equatorial circular photon orbits under perturbations in the z -direction by considering the sign of the quadratic $3d^2 - 4R^2$. Using the fact that stable orbits satisfy $P(R) = 0$, we deduce that $d^2 \left(d - \sqrt{\frac{16}{27}} \right) > 0$. The inner orbit is therefore stable under perturbations out of the plane if $d > d_1 = \sqrt{\frac{16}{27}}$ (and it exists if $d < d_2$).

Summarising these results, we have demonstrated that stable equatorial photon orbits exist for equal-mass Majumdar–Papapetrou di-holes with coordinate separations in the range $d_1 < d < d_2$. We remark that the upper bound of this result ($d < d_2$) is consistent with Coelho and Herdeiro’s results [223]; and with those of Wünsch *et al.* [288].

We now focus on the special case $d = 1$. In this case, the polynomial $P(R)$ can be factorised as $P(R) = (R - 1)(R^2 - R - 1)$. The (positive) roots are therefore $R = 1$ ($\rho = \frac{\sqrt{3}}{2}$) and $R = \varphi$ ($\rho = \sqrt{\varphi^2 - \frac{1}{4}} = \frac{1}{2}5^{1/4}\varphi^{3/2}$), where $\varphi = \frac{1}{2}(1 + \sqrt{5})$ is the golden ratio. Under radial perturbations, the inner (outer) orbit at $R = 1$ ($R = \varphi$) is stable (unstable).

In Section 6.3.2, we derive a condition for the existence of stable photon orbits in stationary axisymmetric electrovacuum spacetimes. We may employ this as a consistency check for the results derived above in the case of an equal-mass Majumdar–Papapetrou di-hole. A key result of Section 6.3.2 is that (6.32) must be satisfied at stationary points, i.e., stationary points of h must satisfy

$$\text{tr } \mathcal{H}(h) = -\frac{2}{\rho} |\mathbf{W}|^2, \quad \mathbf{W} = h\nabla A_t - \nabla A_\phi, \quad (\text{B.5})$$

where $A_t = \frac{1}{U}$ and $A_\phi = 0$ for the Majumdar–Papapetrou solution. Introducing the radial coordinate $R = \sqrt{\rho^2 + \frac{d^2}{4}}$, we find that the existence condition (B.5) is satisfied if $P(R) = R^3 - 2R^2 + d^2 = 0$, which is equivalent to condition (B.2). At a stationary point, the black hole separation can therefore be expressed in terms of the radial coordinate R as $d^2 = 2R^2 - R^3$.

In the case of an equal-mass di-hole, there is an equatorial symmetry. The stability condition for equatorial circular orbits derived in Section 6.3.2 therefore reduces to $0 <$

$-h_{,\rho\rho} < \frac{2}{\rho} |\mathbf{W}|^2$. Rewriting this condition in terms of the radial coordinate R , and making the replacement $d^2 = 2R^2 - R^3$, we arrive at the inequality $\frac{2}{3} < R < \frac{4}{3}$. Noting that $d = \sqrt{2R^2 - R^3}$ is monotonically increasing on the domain $R \in [\frac{2}{3}, \frac{4}{3}]$, the inequality on R may be recast in terms of the separation as $\sqrt{\frac{16}{27}} < d < \sqrt{\frac{32}{27}}$. Hence, we have recovered the existence condition for stable equatorial circular null geodesics in the equal-mass Majumdar–Papapetrou di-hole spacetime.

B.2 Non-planar circular photon orbits

Locating the stationary points of h which do not lie in the equatorial plane poses more of a challenge; however, it is possible to make progress through the introduction of elliptic coordinates (ξ, η) , which are related to (ρ, z) via

$$\rho = \frac{d}{2} \sinh \xi \sin \eta, \quad z = \frac{d}{2} \cosh \xi \cos \eta. \quad (\text{B.6})$$

Making the further substitutions $X = \cosh \xi$ and $Y = \cos \eta$, (B.6) may be recast as

$$\rho = \frac{d}{2} \sqrt{X^2 - 1} \sqrt{1 - Y^2}, \quad z = \frac{d}{2} XY. \quad (\text{B.7})$$

To ensure that $\rho > 0$, we must have $X \in (1, \infty)$ and $Y \in (-1, 1)$. In these coordinates, we have

$$U = 1 + \frac{4X}{d(X^2 - Y^2)}, \quad (\text{B.8})$$

and the effective potential is

$$h = \frac{d}{2} \sqrt{X^2 - 1} \sqrt{1 - Y^2} \left(1 + \frac{4X}{d(X^2 - Y^2)} \right)^2, \quad (\text{B.9})$$

Stationary points satisfy the pair of equations $h_{,X} = 0 = h_{,Y}$. It is quick to establish that these conditions yield the system of equations

$$dX^5 - 2dX^3Y^2 + dXY^4 + 4X^4 + 12X^2Y^2 - 16X^2 = 0, \quad (\text{B.10})$$

$$dX^5 - 2dX^3Y^2 + dXY^4 - 4X^4 - 12X^2Y^2 + 8X^2 + 8Y^2 = 0. \quad (\text{B.11})$$

Subtracting (B.11) from (B.10) and simplifying yields

$$X^4 + 3X^2Y^2 - 3X^2 - Y^2 = 0. \quad (\text{B.12})$$

This relationship allows us to express $Y^2 = \cos^2 \eta$ as a function of $X = \cosh \xi$:

$$Y^2 = \frac{X^2(X^2 - 3)}{1 - 3X^2}. \quad (\text{B.13})$$

Remarkably, this relationship is independent of the value of the black hole separation d . Now, addition of (B.11) and (B.10) results in

$$dX^5 - 2dX^3Y^2 + dXY^4 - 4X^2 + 4Y^2 = 0. \quad (\text{B.14})$$

Eliminating Y from (B.14) using (B.13) leads to a quintic in the X -coordinate:

$$dX^5 - dX^3 - 3X^2 + 1 = 0. \quad (\text{B.15})$$

For general values of the separation d , one is not able to factorise the left-hand side of (B.15), or to find closed-form expressions for its roots in terms of d . However, when $d = 1$, the quintic factorises as $X^5 - X^3 - 3X^2 + 1 = (X^3 + X^2 + X - 1)(X^2 - X - 1)$. The cubic factor has no roots with $X > 1$ (i.e., $\rho > 0$). The quadratic factor has one such root at $X = \varphi$. Using the relationship (B.13), $X = \varphi$ implies that $Y = \pm \frac{1}{\varphi^2}$. In the case $d = 1$, the stationary point conditions are satisfied by $\cosh \xi = \varphi$, $\cos \eta = \pm \frac{1}{\varphi^2}$, i.e., $(\rho, z) = (\frac{1}{2}5^{1/4}\varphi^{-1/2}, \pm \frac{1}{2\varphi})$. Inserting these values into the effective potential $h = \rho U^2$, it is quick to verify that these two non-planar saddle points lie on the same contour as the equatorial saddle point, given by $h = p_\phi = \frac{1}{2}5^{5/4}\varphi^{3/2}$.

The linearity of (B.15) in d allows us to express the separation in terms of the X -coordinate as

$$d = \frac{3X^2 - 1}{X^3(X^2 - 1)}. \quad (\text{B.16})$$

Therefore, it is possible to substitute a closed-form value of X into the right-hand side of (B.16) to find the corresponding value of d . Recalling that we require $0 < Y^2 < 1$, the relation (B.13) may then be used as a consistency check. For example, $X = \sqrt{2}$ is a solution to (B.15) if $d = \frac{5\sqrt{2}}{4}$. Using (B.13), it is quick to check that $Y^2 = \frac{2}{5}$, as required.

Consider (B.16) for a general value of d . The conditions $X > 1$ and $0 < Y^2 < 1$ together imply that we require $X^2 < 3$ for a solution. The critical value $X = \sqrt{3}$ corresponds to a separation of $d = d_1$. For $d < d_1$, there are no circular photon orbits out of the equatorial plane.

Finally, let us consider the limit $X \rightarrow 1$ (i.e., $\rho \rightarrow 0$). Performing a first order perturbative expansion with $X \sim 1 + \varepsilon$, where $\varepsilon \ll 1$, we have $X^2 \sim 1 + 2\varepsilon$. Then, through (B.13),

$$Y^2 = \frac{(1 + 2\varepsilon)(2 - 2\varepsilon)}{2 + 6\varepsilon} \sim 1 - 2\varepsilon. \quad (\text{B.17})$$

Hence, $Y^2 \sim X^{-2}$ as $X \rightarrow 1$. For widely separated black holes with $d \gg 1$, $X \sim 1 + \frac{1}{d}$ and $Y \sim 1 - \frac{1}{d}$. In this case, the equal-mass Majumdar–Papapetrou di-hole will resemble a pair of isolated extremal Reissner–Nordström black holes, with unstable circular photon orbits around each of the sources at $\rho \sim 1$ and $z \sim \pm \frac{d}{2}$. (For a single Reissner–Nordström black hole of mass $M = 1$, there is an unstable equatorial orbit at $\rho = 1$.)

Appendix C

Hénon–Heiles Hamiltonian system

In the early 1960s, Hénon and Heiles [117] investigated the non-linear motion of stars about the galactic centre. The main purpose of this work was to obtain a third integral of motion – in addition to the energy and angular momentum – in galactic dynamics. To do this, Hénon and Heiles consider a simple two-dimensional Hamiltonian, described in phase space coordinates $\{x, y, p_x, p_y\}$ by

$$H = \frac{1}{2} (p_x^2 + p_y^2) + V(x, y), \quad (\text{C.1})$$

$$V = \frac{1}{2}k (x^2 + y^2) + \lambda \left(x^2y - \frac{1}{3}y^3 \right), \quad (\text{C.2})$$

where $k, \lambda \in \mathbb{R}$ are parameters. If λ is small, the potential resembles that of a two-dimensional isotropic harmonic oscillator, with two cubic perturbation terms.

The system exhibits a threefold rotational symmetry in the (x, y) -plane. This is seen most easily by transforming to planar polar coordinates $x = r \cos \theta$, $y = r \sin \theta$, in which the potential reads $V = \frac{1}{2}kr^2 + \frac{1}{3}\lambda r^3 \sin 3\theta$. The invariance of the potential under transformations of the form $\theta \mapsto \theta + \frac{2n\pi}{3}$ ($n \in \mathbb{Z}$) is now manifest.

For simplicity, we consider the case $k = 1 = \lambda$. The Hamiltonian H is a constant of the motion, say E , which corresponds to the total energy of the system. The equations of motion, i.e., Hamilton's equations, read

$$\dot{x} = p_x, \quad \dot{y} = p_y, \quad \dot{p}_x = -x - 2xy, \quad \dot{p}_y = -y - x^2 + y^2, \quad (\text{C.3})$$

where an overdot denotes differentiation with respect to some time parameter, say t . The coupled non-linear system (C.3) is not Liouville integrable, and has no closed-form solution. One should therefore anticipate rich dynamics.

The full phase space of the Hénon–Heiles system is described by two pairs of conjugate variables $\{x, y, p_x, p_y\}$, and an energy constraint $H = E$. For a fixed value of the energy E , we must have

$$\frac{1}{2} (p_x^2 + p_y^2) \leq E, \quad (\text{C.4})$$

$$V(x, y) \leq E. \quad (\text{C.5})$$

The inequality (C.4) indicates that the kinetic energy of the particle is bounded above; the inequality (C.5) demarcates the allowed regions of configuration space which are accessible by the particle.

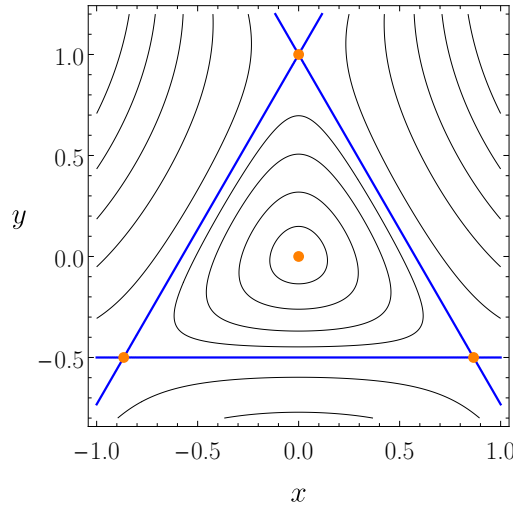


Figure C.1. Equipotential curves of the two-dimensional Hénon–Heiles potential for a selection of values of the conserved energy E . The stationary points of V (three saddle points and a local minimum) are represented by orange points. The local minimum is at the origin $(0, 0)$, and has $V = 0$. The three saddles are located at $(0, 1)$ and $(\pm\frac{\sqrt{3}}{2}, -\frac{1}{2})$, and are connected by the critical contour $V = E^*$. For energies with $0 < E \leq E^*$ the contours form closed curves; whereas for $E > E^*$ the contours are open, and a particle can access spatial infinity through one of three escapes.

Orbits of constant x and y satisfy $\nabla V = (V_{,x}, V_{,y}) = (0, 0)$, with $V = E$; in other words, such orbits are stationary points of the potential. Locating these stationary points is tantamount to solving the system (C.3) for x and y , with $p_x = 0 = p_y$. A stationary point is stable if it is a local minimum of the potential V , i.e., if $\det \mathcal{H}(V) > 0$ and $\text{tr} \mathcal{H}(V) > 0$, where $\mathcal{H}(V)$ denotes the Hessian matrix of the potential V . The Hénon–Heiles potential admits a local minimum $(x, y) = (0, 0)$, where $V = 0$. Moreover, there are three saddle points which lie on the vertices of an equilateral triangle with coordinates $(0, 1)$ and $(\pm\frac{\sqrt{3}}{2}, -\frac{1}{2})$. The three saddles are connected by a single contour $V = E^* = \frac{1}{6}$ which encloses the local minimum.

In Figure C.1, we show the equipotential curves for the Hénon–Heiles potential (C.2) (i.e., curves of constant V) in the (x, y) -plane. For $E > E^*$ ($E \leq E^*$), the Hénon–Heiles equipotential lines are open (closed). The former case is related to the investigation of Wada structures in binary black hole shadows, presented in Chapter 5. The latter case is related to the study of bounded stable photon orbits around Majumdar–Papapetrou di-holes, discussed in Chapter 6. We discuss each case separately below.

C.1 Open system

For $E > E^*$, the Hénon–Heiles system is open: orbits are permitted to escape to spatial infinity. For energies close to the threshold $E \gtrsim E^*$, there are three narrow channels which connect the scattering region to spatial infinity. In this regime, the Hénon–Heiles system is an example of an open Hamiltonian system with three escapes [112].

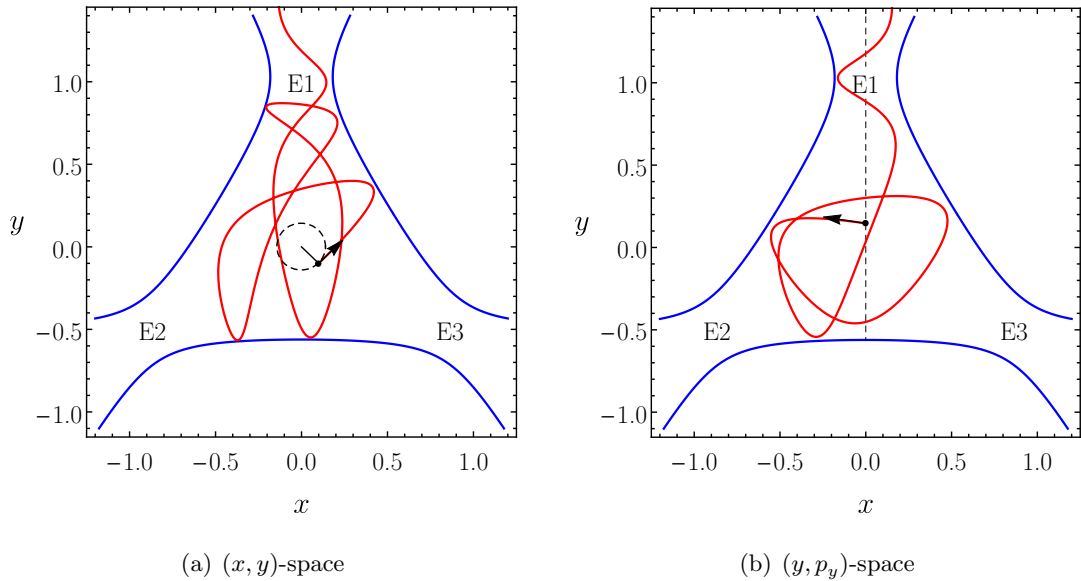


Figure C.2. Initial conditions for two-dimensional exit basins in the Hénon–Heiles open Hamiltonian system in (a) the (x, y) -plane; and (b) the (y, p_y) -plane. In both cases, the contour $V = E^* + \Delta E$ is shown in blue, and a trajectory which escapes the scattering region through Exit 1 is shown in red.

In an open Hamiltonian system with multiple escapes, an exit basin is defined as the set of initial conditions which lead to a certain exit in phase space (see Section 2.6.5). As described above, the Hénon–Heiles system is a two-dimensional time-independent Hamiltonian system: the phase space is described by two pairs of conjugate variables $\{x, y, p_x, p_y\}$, and an energy constraint $H = E$. One must fix three initial conditions to define a trajectory: the “physical” phase space is therefore three-dimensional. In order to visualise the exit basins of such a Hamiltonian system, one can employ a surface of section (or Poincaré section) to reduce the dimensionality of the phase space by one; see Section 2.6.5 for a review. In [246], the authors employ two choices of initial conditions which define a two-dimensional surface of section; these initial conditions are chosen so that the exit basins contain a *Lyapunov orbit* – an unstable periodic orbit which exists for all energies above the threshold value, such that any trajectory which crosses the Lyapunov orbit with its velocity vector pointing outwards is forced to escape without returning to the scattering region.

The first choice of initial conditions is defined as follows. For each pair of coordinates (x, y) which satisfy the inequality (C.5) (i.e., which lie in the allowed region of configuration space), we choose the initial momentum (p_x, p_y) to be tangent (in the anti-clockwise sense) to the circle centred on $(0, 0)$ which passes through the initial condition (x, y) . In other words, the two-dimensional surface of section is defined by the constraint $xp_x + yp_y = 0$, with $xp_y - yp_x > 0$. (See Figure C.2(a) for the set-up of these initial conditions.)

For our second choice of initial conditions, we fix $x = 0$ and vary the initial values of y and p_y ; the initial value of the momentum p_x is determined by rearranging the Hamiltonian constraint $H = E$. (The set-up of these initial conditions is shown in Figure C.2(b).)

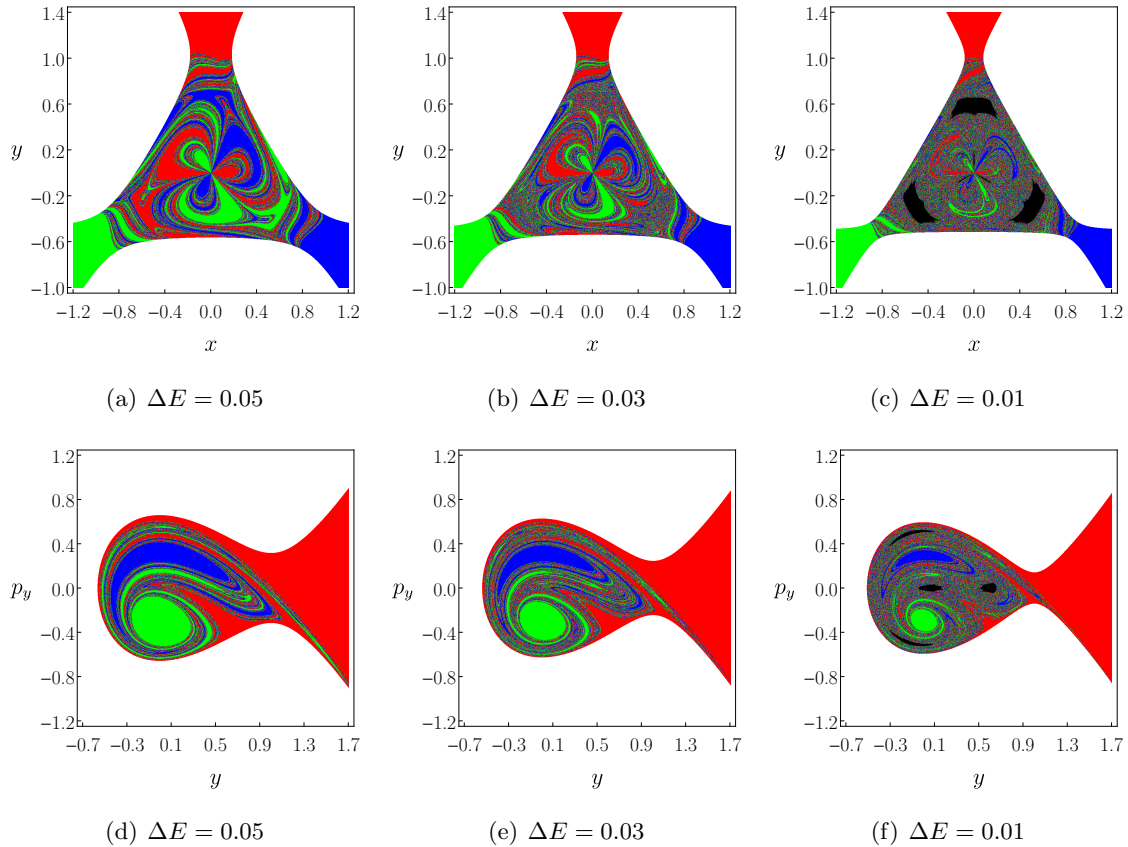


Figure C.3. Exit basins in the Hénon–Heiles open Hamiltonian system for (a)–(c) initial conditions in the (x, y) -plane; and (d)–(f) initial conditions in (y, p_y) -plane. The basins corresponding to Exits 1, 2, and 3 are plotted in red, green, and blue, respectively.

In Figure C.3, we show the exit basins for the Hénon–Heiles system for a selection of values of the energy parameter $\Delta E = E - E^*$, using the two types of initial conditions described above. In practice, the exit basins are realised by numerically integrating Hamilton’s equations for a fine grid of initial conditions in either the (x, y) -plane or the (y, p_y) -plane. The initial conditions are coloured red, green or blue according to whether they leave the scattering region through Exit 1, Exit 2 or Exit 3, respectively.

Figure C.3 shows the effect of varying the energy ΔE on the structure of the basins. As $\Delta E \rightarrow 0$ (i.e., as $E \rightarrow E^* = \frac{1}{6}$ from above), the basins become more “fractalised”: the three (disconnected) basins are intertwined in a complicated fashion, and their boundary is highly intricate. Moreover, Kolmogorov–Arnold–Moser islands of stability (depicted in black) remain as “deterministic islands” in a “random sea”. The exit basins of the Hénon–Heiles system have been explored in detail by Aguirre *et al.* [246], and in the limit of small escapes by Aguirre and Sanjuán [227].

The exit basins of the Hénon–Heiles system possess the Wada property: any point on the boundary of one basin is on the boundary of the other two basins (see Section 5.5). This was verified numerically using the Nusse–Yorke method by Aguirre *et al.* in 2001 [246]. Since then, a range of numerical algorithms have been proposed to test for the Wada property in the exit basins of open Hamiltonian systems, using the Hénon–

Heiles system as a standard example; these include the grid method [257], the merging method [248], and the saddle–straddle method [321]. In practice, the existence of Wada basins in phase space mean that the system exhibits extreme sensitive dependence on initial conditions: a small uncertainty in fixing the initial conditions in the neighbourhood of a Wada boundary results in high levels of indeterminism in the final state of the system, despite the underlying dynamics being fully deterministic. This is discussed in more depth throughout Chapter 5.

Additional studies of the fractal structures which arise in the exit basins of the Hénon–Heiles system have been conducted by Barrio *et al.* [322] and Zotos [323].

C.2 Closed system

For conserved energies in the range $0 < E \leq \frac{1}{6}$, the Hénon–Heiles system is closed: orbits are kinematically bounded by the energy, and the motion is restricted to an open subset of the (x, y) -plane. The inequality (C.4) indicates that the kinetic energy of the particle is bounded above; the inequality (C.5) demarcates the allowed regions of configuration space which are accessible by the particle. Any trajectory which starts inside a closed contour $V = E$ will remain so for all $t > 0$.

In order to visualise the complex nature of trajectories in phase space, one may use a Poincaré section (see Section 2.6.5). A rearrangement of the Hamiltonian constraint $H = E$ allows us to express one of the phase space coordinates, say p_x , in terms of the other three. The dimensionality of the “physical” phase space is therefore reduced by one, and the motion is confined to a three-dimensional energy hypersurface. Since the motion is kinematically bounded, the orbits will repeatedly intersect any embedded two-dimensional surface, e.g. the (y, p_y) -plane ($x = 0$). If there exists an additional conserved quantity involving y and p_y , then it could be used to express $p_y = p_y(y)$; the points of intersection between the trajectory and the surface $x = 0$ will therefore lie on a smooth curve.

To illustrate the utility of the Poincaré section, let us consider the dynamics of bounded trajectories the Hénon–Heiles system for a selection of energies $0 < E \leq \frac{1}{6}$, where the surface of section is taken to be $x = 0$, i.e., the (y, p_y) -plane (see Section 2.6.5). The Hamiltonian constraint $H = E$, coupled with the fact that $p_x^2 \geq 0$ imply that motion is restricted to a subset of the (y, p_y) -plane given by

$$\frac{1}{2}p_y^2 + \frac{1}{2}y^2 - \frac{1}{3}y^3 \leq E. \quad (\text{C.6})$$

The extreme values of p_y occur where $y = 0$, i.e., $p_y = \pm\sqrt{2E}$. Similarly, the extreme values of y occur where $p_y = 0$, i.e., at the two real roots of the cubic $Q(y) = \frac{1}{3}y^3 - \frac{1}{2}y^2 + E$ which satisfy the inequality (C.5) with $x = 0$. The cubic $Q(y)$ has discriminant $\Delta_y(Q) = \frac{1}{2}E(1 - 6E)$; two such roots are therefore guaranteed to exist for energies in the range $0 < E \leq \frac{1}{6}$.

On the surface $x = 0$, we can express p_x as a function of y and p_y as

$$p_x = \pm\sqrt{2E - p_y^2 - y^2 + \frac{2}{3}y^3}. \quad (\text{C.7})$$

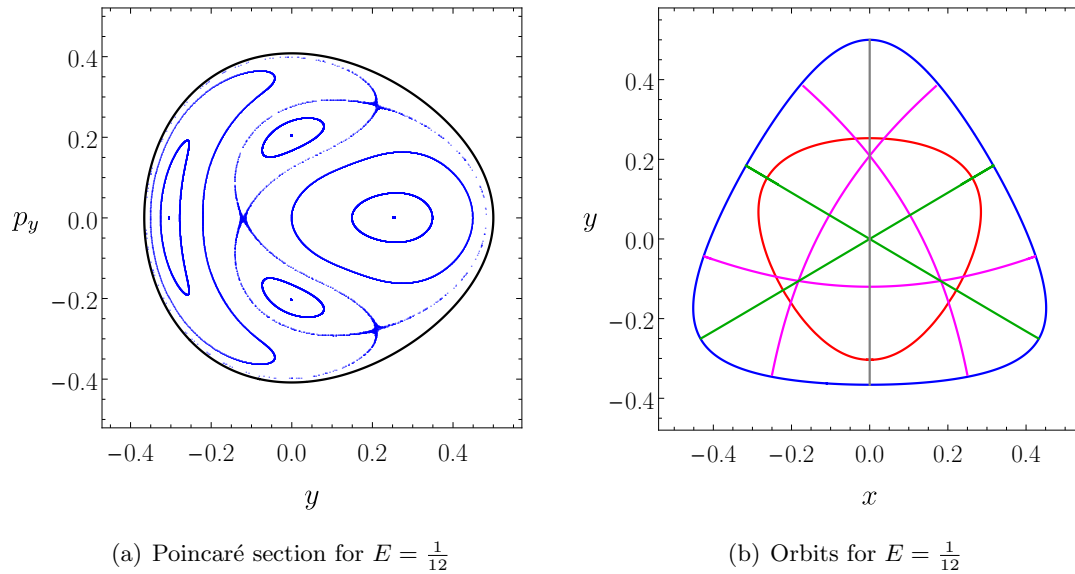


Figure C.4. (a) Poincaré section in the (y, p_y) -plane for the Hénon–Heiles Hamiltonian system with conserved total energy $E = \frac{1}{12}$. (b) A selection of (stable and unstable) periodic orbits which lie within the closed contour $V = E = \frac{1}{12}$ [blue]: a rotational 1 : 1 resonant orbit (or “loop orbit”) [red]; a pair of librational 1 : 1 resonant orbits (or “linear orbits”) [green]; and three unstable linear orbits [magenta].

In practice, the construction of a Poincaré section involves numerically integrating trajectories in the three-dimensional phase space for a selection of initial conditions, recording the values of y and p_y each time the trajectory intersects $x = 0$ with $p_x > 0$, i.e., taking the upper (+) sign in (C.7).

Figure C.4(a) shows an example of a Poincaré section for the Hénon–Heiles system with $E = \frac{1}{12}$. There are four regions which contain closed oval-shaped curves, surrounding four (stable) elliptic fixed points. These four regions are separated by a continuous curve, which has self-intersections at three (unstable) hyperbolic fixed points. (The Poincaré section is clearly symmetric about the y -axis, due to the fact that the Hamiltonian (C.1) is invariant under the transformation $p_y \mapsto -p_y$.) In Figure C.4(b), we present a selection of periodic orbits, bounded by the closed contour $V = E = \frac{1}{12}$, shown in blue. The red curve is a “rotational” orbit, corresponding to the pair of elliptic fixed points in the plane $p_y = 0$. The green curves are “librational” orbits, corresponding to the non-planar elliptic fixed points. The rotational and librational orbits are partitioned by a “separatrix”, which intersects itself at the three elliptic fixed points; the corresponding unstable periodic orbits are shown as magenta curves.

We now consider the effect of increasing the energy E , towards the threshold $E = \frac{1}{6}$. In Figure C.5, we present a selection of Poincaré sections and orbits for $E \in \{\frac{1}{8}, \frac{1}{6}\}$.

Figure C.5(a) shows that, when $E = \frac{1}{8}$, there exist regions containing closed oval-shaped curves which surround elliptic fixed points. However, these regular “islands” are no longer demarcated by a smooth separatrix: the intermediate regions are filled with irregular orbits. In fact, all of the points in the “speckled” region of Figure C.5(a) are

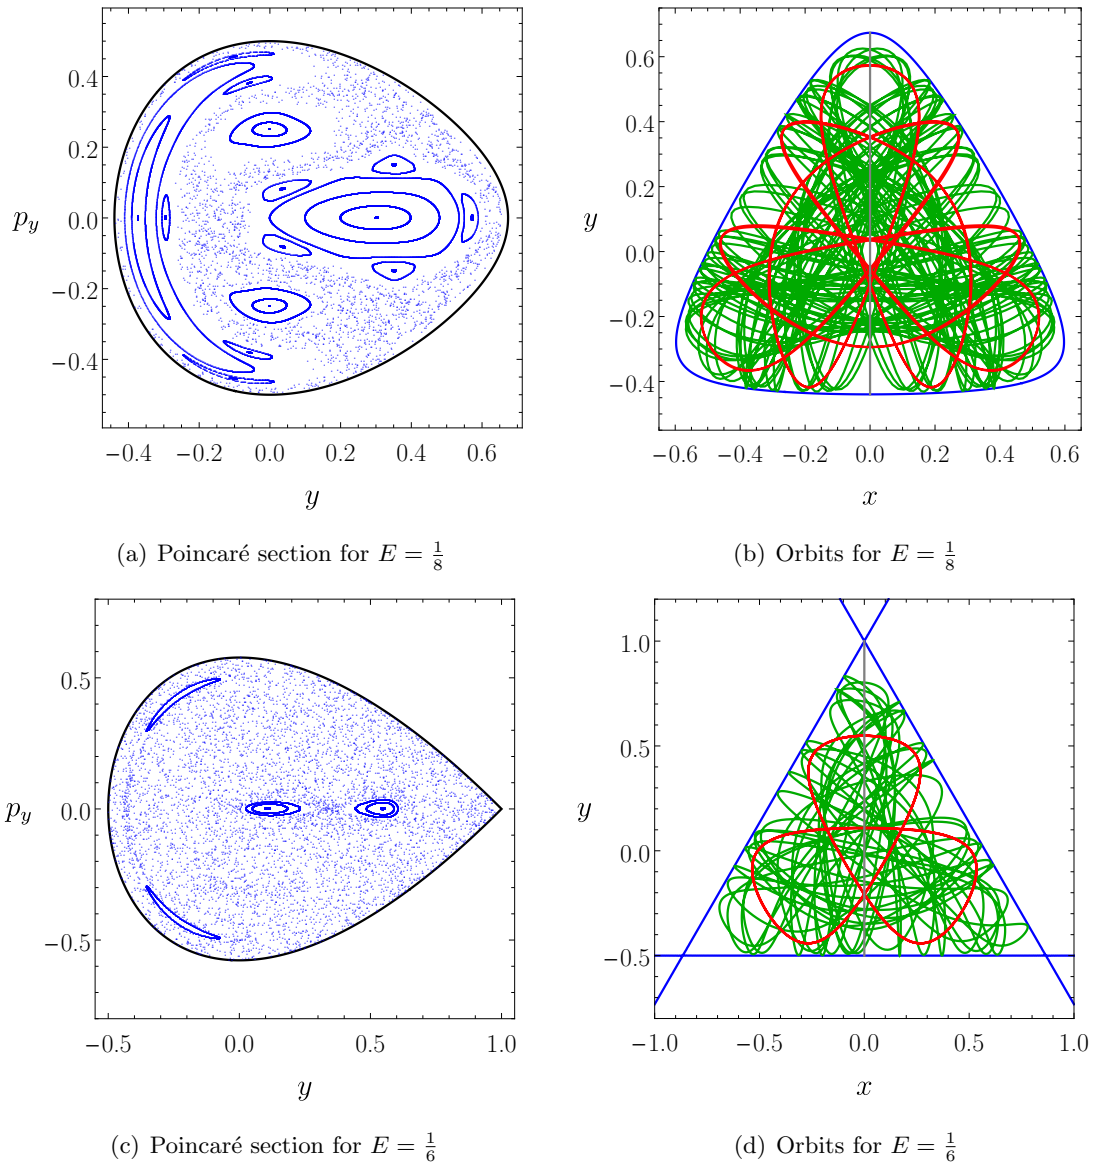


Figure C.5. (a) Poincaré section in the (y, p_y) -plane for $E = \frac{1}{8}$. We observe regions of regular motion around elliptic fixed points, and chaotic motion. (b) Orbits bounded by the closed contour $V = E = \frac{1}{8}$ [blue]: a 5 : 5 resonant orbit [red]; and a chaotic orbit [green]. (c) Poincaré section for $E = \frac{1}{6}$ with chaos dominant. (d) Orbits which lie within the closed contour $V = E = \frac{1}{6}$ [blue]: a 2 : 2 resonant orbit [red]; and a chaotic orbit [green], which wanders through the strange attractor region of phase space.

generated by the repeated intersection of a single orbit with the surface of section. In Figure C.5(b), we present a selection of orbits for $E = \frac{1}{8}$. A resonant periodic orbit is shown in red; and an irregular (chaotic) orbit is shown in green.

Increasing the energy to its maximum value $E = \frac{1}{6}$, we see that chaos becomes dominant. Figure C.5(c) shows the Poincaré section for $E = \frac{1}{6}$. We see that, in this case, chaotic orbits fill almost all of the available phase space; however, there exist very small regular islands, which surround elliptic fixed points. We show a selection of orbits for $E = \frac{1}{6}$ in Figure C.5(d). The stable periodic orbit – corresponding to the elliptic fixed point of the Poincaré map – is shown in red. We also show an example of a chaotic orbit in green.

More comprehensive discussions of bounded motion in the Hénon–Heiles Hamiltonian system as well as a detailed description of the Poincaré sections can be found in [103, 296, 324], for example.

Appendix D

Bretón–Manko–Aguilar di-hole solution

The $N = 2$ Bretón–Manko–Aguilar [232] solution describes a pair of aligned Reissner–Nordström particles in static equilibrium. The geometry is a solution to the Einstein–Maxwell equations, and is described by the static Weyl line element

$$ds^2 = -f dt^2 + \frac{1}{f} [e^{2\gamma} (d\rho^2 + dz^2) + \rho^2 d\phi^2], \quad (\text{D.1})$$

where $f = f(\rho, z)$, $\gamma = \gamma(\rho, z)$, and the only non-zero component of the electromagnetic gauge potential is $A_t(\rho, z)$.

Manko [295] presents the general five-parameter Bretón–Manko–Aguilar solution in terms of the “physical parameters”: the black hole masses M_{\pm} ; the black hole charges Q_{\pm} ; and the separation between the centres d . In the so-called physical parametrisation, the metric functions and electrostatic potential take the form

$$f = \frac{A^2 - B^2 + C^2}{(A + B)^2}, \quad e^{2\gamma} = \frac{A^2 - B^2 + C^2}{16\sigma_+^2 \sigma_-^2 (\nu + 2\kappa)^2 r_1 r_2 r_3 r_4}, \quad A_t = \frac{C}{A + B}, \quad (\text{D.2})$$

where the functions A , B and C are given by

$$A = \sigma_+ \sigma_- [\nu (r_1 + r_2) (r_3 + r_4) + 4\kappa (r_1 r_2 + r_3 r_4)] - (\mu^2 \nu - 2\kappa^2) (r_1 - r_2) (r_3 - r_4), \quad (\text{D.3})$$

$$B = 2\sigma_+ \sigma_- [(\nu M_+ + 2\kappa M_-) (r_1 + r_2) + (\nu M_- + 2\kappa M_+) (r_3 + r_4)] - 2\sigma_+ [\nu \mu (Q_- + \mu) + 2\kappa (dM_- + \mu Q_+ - \mu^2)] (r_1 - r_2) - 2\sigma_- [\nu \mu (Q_+ - \mu) - 2\kappa (dM_+ - \mu Q_- - \mu^2)] (r_3 - r_4), \quad (\text{D.4})$$

$$C = 2\sigma_+ \sigma_- \{[\nu (Q_+ - \mu) + 2\kappa (Q_- + \mu)] (r_1 + r_2) + [\nu (Q_- + \mu) + 2\kappa (Q_+ - \mu)] (r_3 + r_4)\} - 2\sigma_+ (\mu \nu M_- + 2\kappa (\mu M_+ + dQ_- + \mu d)) (r_1 - r_2) - 2\sigma_- [\mu \nu M_+ + 2\kappa (\mu M_- - dQ_+ + \mu d)] (r_3 - r_4), \quad (\text{D.5})$$

with

$$\nu = d^2 - \sigma_+^2 - \sigma_-^2 + 2\mu^2, \quad \kappa = M_+M_- - (Q_+ - \mu)(Q_- + \mu), \quad (\text{D.6})$$

$$\mu = \frac{M_-Q_+ - M_+Q_-}{M_+ + M_- + d}, \quad \sigma_{\pm} = \sqrt{M_{\pm}^2 - Q_{\pm}^2 \pm 2\mu Q_{\pm}}, \quad (\text{D.7})$$

$$r_1 = \sqrt{\rho^2 + \left(z - \frac{d}{2} - \sigma_-\right)^2}, \quad r_2 = \sqrt{\rho^2 + \left(z - \frac{d}{2} + \sigma_-\right)^2}, \quad (\text{D.8})$$

$$r_3 = \sqrt{\rho^2 + \left(z + \frac{d}{2} - \sigma_+\right)^2}, \quad r_4 = \sqrt{\rho^2 + \left(z + \frac{d}{2} + \sigma_+\right)^2}. \quad (\text{D.9})$$

In [295], Manko derives a concise formula for the force imparted by the Weyl strut which holds the sources in static equilibrium:

$$\mathcal{F} = \frac{\kappa}{\nu - 2\kappa} = \frac{M_+M_- - (Q_+ - \mu)(Q_- + \mu)}{d^2 - (M_+ + M_-)^2 + (Q_+ + Q_-)^2}. \quad (\text{D.10})$$

Appendix E

Tetrad calculations

E.1 Identities for Killing objects and tetrads

Using the properties of the Killing objects from Chapter 7, we prove the identities (7.55) and (7.56) for the self-dual bivector $\mathcal{F}_{ab} = f_{ab} - ih_{ab}$, where f_{ab} is the Killing–Yano tensor and h_{ab} is the closed conformal Killing–Yano tensor.

Firstly, recall that $f_{ac}h_b{}^c = h_{ac}f_b{}^c = ar \cos \theta g_{ab}$, and $K_{ab} - Q_{ab} = (r^2 - a^2 \cos^2 \theta) g_{ab}$. The “square” of \mathcal{F}_{ab} is then

$$\mathcal{F}_{ac}\mathcal{F}_b{}^c = (f_{ac} - ih_{ac})(f_b{}^c - ih_b{}^c) \quad (\text{E.1})$$

$$= f_{ac}f_b{}^c - h_{ac}h_b{}^c - i(f_{ac}h_b{}^c + h_{ac}f_b{}^c) \quad (\text{E.2})$$

$$= K_{ab} - Q_{ab} - 2iar \cos \theta g_{ab} \quad (\text{E.3})$$

$$= (r^2 - a^2 \cos^2 \theta - 2iar \cos \theta) g_{ab} \quad (\text{E.4})$$

$$= (r - ia \cos \theta)^2 g_{ab}. \quad (\text{E.5})$$

This proves (7.55). Similarly, the product of \mathcal{F}_{ab} with its complex conjugate is

$$\mathcal{F}_{ac}\overline{\mathcal{F}}_b{}^c = (f_{ac} - ih_{ac})(f_b{}^c + ih_b{}^c) \quad (\text{E.6})$$

$$= f_{ac}f_b{}^c + h_{ac}h_b{}^c + i(f_{ac}h_b{}^c - h_{ac}f_b{}^c) \quad (\text{E.7})$$

$$= K_{ab} + Q_{ab}. \quad (\text{E.8})$$

This proves result (7.56).

We now derive the transport equation (7.61) by applying the derivative operator $\tilde{D} = \tilde{k}^a \nabla_a$ to the tetrad legs \tilde{m}^a and \tilde{n}^a , which are defined in terms of the self-dual bivector \mathcal{F}_{ab} as

$$\tilde{m}^a = \frac{1}{\sqrt{2K}} \mathcal{F}^a{}_b \tilde{k}^b, \quad \tilde{n}^a = \frac{1}{2K} \mathcal{F}^{ab} \overline{\mathcal{F}}_{bc} \tilde{k}^c. \quad (\text{E.9})$$

First, recall that $f_{ab;c} = f_{[ab;c]}$, and $h_{ab;c} = g_{ac}\xi_b - g_{cb}\xi_a$, where $\xi^a = \xi_{(t)}^a$ is the primary Killing vector field. We also note that $\tilde{k}^a \xi_a = -E$, where E is the photon’s conserved

energy. Using these results, the transport equation for \tilde{m}^a is then

$$\tilde{D}\tilde{m}_a = \frac{1}{\sqrt{2K}}\tilde{k}^c \left(\mathcal{F}_{ab}\tilde{k}^b \right)_{;c} \quad (\text{E.10})$$

$$= \frac{1}{\sqrt{2K}}\mathcal{F}_{ab;c}\tilde{k}^b\tilde{k}^c \quad (\text{E.11})$$

$$= \frac{1}{\sqrt{2K}}(f_{ab;c} - ih_{ab;c})\tilde{k}^b\tilde{k}^c \quad (\text{E.12})$$

$$= \frac{1}{\sqrt{2K}} \left[f_{[ab;c]} - i(g_{ac}\xi_b - g_{cb}\xi_a) \right] \tilde{k}^b\tilde{k}^c \quad (\text{E.13})$$

$$= -\frac{i}{\sqrt{2K}} \left(\tilde{k}^b\xi_b \right) \tilde{k}_a \quad (\text{E.14})$$

$$= \frac{iE}{\sqrt{2K}}\tilde{k}_a. \quad (\text{E.15})$$

This is precisely (7.61).

E.2 Relationship with Marck's tetrads

In Section 7.4 we claim that the complex null tetrad $\{\tilde{k}^a, \tilde{n}^a, \tilde{m}^a, \overline{\tilde{m}^a}\}$ constructed from the symmetries of Kerr spacetime is equivalent to that constructed using the legs of Marck's quasi-orthonormal tetrad. In particular,

$$\tilde{n}^a = \frac{1}{2K}\mathcal{F}^{ab}\overline{\mathcal{F}}_{bc}\tilde{k}^c = -\tilde{\lambda}_3^a, \quad \tilde{m}^a = \frac{1}{\sqrt{2K}}\mathcal{F}^a_b\tilde{k}^b = \frac{1}{\sqrt{2}}(\tilde{\lambda}_2^a + i\tilde{\lambda}_1^a). \quad (\text{E.16})$$

This is best demonstrated by projecting onto Carter's tetrad as follows. First, we note that $\tilde{k}^a = \tilde{\lambda}_0^a$ by definition. For notational simplicity, we introduce the coordinate $y = a \cos \theta$.¹ When projected onto Carter's symmetric tetrad, the Killing objects take the form [50]

$$f_{(\alpha)(\beta)} = \begin{bmatrix} 0 & -y & 0 & 0 \\ y & 0 & 0 & 0 \\ 0 & 0 & 0 & -r \\ 0 & 0 & r & 0 \end{bmatrix}, \quad h_{(\alpha)(\beta)} = \begin{bmatrix} 0 & r & 0 & 0 \\ -r & 0 & 0 & 0 \\ 0 & 0 & 0 & -y \\ 0 & 0 & y & 0 \end{bmatrix}, \quad (\text{E.17})$$

$$K_{(\alpha)(\beta)} = \begin{bmatrix} y^2 & 0 & 0 & 0 \\ 0 & -y^2 & 0 & 0 \\ 0 & 0 & r^2 & 0 \\ 0 & 0 & 0 & r^2 \end{bmatrix}, \quad Q_{(\alpha)(\beta)} = \begin{bmatrix} r^2 & 0 & 0 & 0 \\ 0 & -r^2 & 0 & 0 \\ 0 & 0 & y^2 & 0 \\ 0 & 0 & 0 & y^2 \end{bmatrix}. \quad (\text{E.18})$$

Recall that $\Sigma = r^2 + a^2 \cos^2 \theta = r^2 + y^2$. Using the definition (E.16) and the identity $\mathcal{F}_{ac}\overline{\mathcal{F}}_b^c = K_{ab} + Q_{ab}$, we find that the leg n^a , when projected onto Carter's symmetric

¹In fact, $y = a \cos \theta$ is one of the so-called *canonical coordinates* [50].

tetrad, takes the form

$$\tilde{n}^{(\alpha)} = -\frac{1}{2K} \left(K^{(\alpha)}_{(\beta)} + Q^{(\alpha)}_{(\beta)} \right) \tilde{\lambda}_0^{(\beta)} \quad (\text{E.19})$$

$$= -\frac{\Sigma}{2K} \begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \tilde{\lambda}_0^{(0)} \\ \tilde{\lambda}_0^{(1)} \\ \tilde{\lambda}_0^{(2)} \\ \tilde{\lambda}_0^{(3)} \end{bmatrix} \quad (\text{E.20})$$

$$= -\frac{\Sigma}{2K} \begin{bmatrix} -\tilde{\lambda}_0^{(0)} \\ -\tilde{\lambda}_0^{(1)} \\ \tilde{\lambda}_0^{(2)} \\ \tilde{\lambda}_0^{(3)} \end{bmatrix}. \quad (\text{E.21})$$

By comparing (E.21) with the components of $\tilde{\lambda}_3^{(\alpha)}$ given in Section 7.4.2, we see that $n^a = -\tilde{\lambda}_3^a$, as claimed.

Similarly, when projected onto Carter's tetrad, we find that the leg \tilde{m}^a may be expressed in the form

$$\tilde{m}^{(\alpha)} = \frac{1}{\sqrt{2K}} \left(f^{(\alpha)}_{(\beta)} - ih^{(\alpha)}_{(\beta)} \right) \tilde{\lambda}_0^{(\beta)} \quad (\text{E.22})$$

$$= \frac{1}{\sqrt{2K}} \left(\begin{bmatrix} 0 & y & 0 & 0 \\ -y & 0 & 0 & 0 \\ 0 & 0 & 0 & -r \\ 0 & 0 & r & 0 \end{bmatrix} + i \begin{bmatrix} 0 & r & 0 & 0 \\ r & 0 & 0 & 0 \\ 0 & 0 & 0 & y \\ 0 & 0 & -y & 0 \end{bmatrix} \right) \begin{bmatrix} \tilde{\lambda}_0^{(0)} \\ \tilde{\lambda}_0^{(1)} \\ \tilde{\lambda}_0^{(2)} \\ \tilde{\lambda}_0^{(3)} \end{bmatrix} \quad (\text{E.23})$$

$$= \frac{1}{\sqrt{2}} \left(\frac{1}{\sqrt{K}} \begin{bmatrix} y\tilde{\lambda}_0^{(1)} \\ y\tilde{\lambda}_0^{(0)} \\ -r\tilde{\lambda}_0^{(3)} \\ r\tilde{\lambda}_0^{(2)} \end{bmatrix} + \frac{i}{\sqrt{K}} \begin{bmatrix} r\tilde{\lambda}_0^{(1)} \\ r\tilde{\lambda}_0^{(0)} \\ y\tilde{\lambda}_0^{(3)} \\ -y\tilde{\lambda}_0^{(2)} \end{bmatrix} \right). \quad (\text{E.24})$$

Comparison of (E.24) with the legs $\tilde{\lambda}_1^{(\alpha)}$ and $\tilde{\lambda}_2^{(\alpha)}$ of Marck's quasi-orthonormal tetrad (Section 7.4.2) shows that $\tilde{m}^a = \frac{1}{\sqrt{2}} (\tilde{\lambda}_2^a + i\tilde{\lambda}_1^a)$. Considering the real and imaginary parts of \tilde{m}^a separately, we see that

$$\tilde{\lambda}_2^a = \frac{1}{\sqrt{K}} f^a_b k^b, \quad \tilde{\lambda}_1^a = -\frac{1}{\sqrt{K}} h^a_b k^b. \quad (\text{E.25})$$

The first of these identities is unsurprising: Marck [52] defines the leg $\tilde{\lambda}_2^a$ to be the unit spacelike vector which is parallel to the vector $v^a = f^a_b k^b$.

E.3 Relationship with Kinnersley's tetrad

We begin with the well-known *Kinnersley tetrad* $\{k^a, n^a, m^a, \bar{m}^a\}$ [325], which in Boyer–Lindquist coordinates has components

$$k^a \partial_a = \frac{1}{\Delta} [(r^2 + a^2) \partial_t + \Delta \partial_r + a \partial_\phi], \quad (\text{E.26})$$

$$n^a \partial_a = \frac{1}{2\Sigma} [(r^2 + a^2) \partial_t - \Delta \partial_r + a \partial_\phi], \quad (\text{E.27})$$

$$m^a \partial_a = \frac{1}{\sqrt{2}(r + ia \cos \theta)} \left(ia \sin \theta \partial_t + \partial_\theta + \frac{i}{\sin \theta} \partial_\phi \right), \quad (\text{E.28})$$

$$\bar{m}^a \partial_a = \frac{1}{\sqrt{2}(r - ia \cos \theta)} \left(-ia \sin \theta \partial_t + \partial_\theta - \frac{i}{\sin \theta} \partial_\phi \right), \quad (\text{E.29})$$

where $\Sigma(r, \theta) = r^2 + a^2 \cos^2 \theta$ and $\Delta(r) = r^2 - 2Mr + a^2$. All inner products of the Kinnersley tetrad are zero, except $k^a n_a = -1$ and $m^a \bar{m}_a = 1$, i.e., it is a complex null tetrad.

We define the *symmetrised Kinnersley tetrad* as

$$\hat{k}^a = \sqrt{\frac{\Delta}{2\Sigma}} k^a, \quad \hat{n}^a = \sqrt{\frac{2\Sigma}{\Delta}} n^a, \quad \hat{m}^a = m^a, \quad \hat{\bar{m}}^a = \bar{m}^a. \quad (\text{E.30})$$

This is a Lorentz boost in the (k, n) -plane; see the transformation (2.110) in Chapter 2. Clearly, the symmetrised null tetrad satisfies the same inner product relationships as the standard Kinnersley tetrad.

One may express the “tilded” complex null tetrad $\{\tilde{k}^a, \tilde{n}^a, \tilde{m}^a, \tilde{\bar{m}}^a\}$ (see Chapter 7) in terms of the symmetrised Kinnersley tetrad as

$$k^a = A \left(e^{+\gamma} \hat{k}^a + e^{-\gamma} \hat{n}^a + e^{-i\beta} \hat{m}^a + e^{+i\beta} \hat{\bar{m}}^a \right), \quad (\text{E.31})$$

$$\tilde{n}^a = \frac{1}{4A} \left(e^{+\gamma} \hat{k}^a + e^{-\gamma} \hat{n}^a - e^{-i\beta} \hat{m}^a - e^{+i\beta} \hat{\bar{m}}^a \right), \quad (\text{E.32})$$

$$\tilde{m}^a = \frac{ie^{i\alpha}}{2} \left(e^{+\gamma} \hat{k}^a - e^{-\gamma} \hat{n}^a + e^{-i\beta} \hat{m}^a - e^{+i\beta} \hat{\bar{m}}^a \right), \quad (\text{E.33})$$

$$\tilde{\bar{m}}^a = -\frac{ie^{-i\alpha}}{2} \left(e^{+\gamma} \hat{k}^a - e^{-\gamma} \hat{n}^a - e^{-i\beta} \hat{m}^a + e^{+i\beta} \hat{\bar{m}}^a \right), \quad (\text{E.34})$$

where we have made the definitions

$$A = \sqrt{\frac{K}{2\Sigma}}, \quad (\text{E.35})$$

$$e^{\pm\gamma} = \frac{\Sigma}{\sqrt{K\Delta}} \left[\frac{E(r^2 + a^2) - aL_z}{\Sigma} \pm \dot{r} \right], \quad (\text{E.36})$$

$$e^{i\beta} = e^{i\alpha} \sqrt{\frac{Z}{\bar{Z}}}, \quad (\text{E.37})$$

$$e^{i\alpha} = \sqrt{\frac{r - ia \sin \theta}{r + ia \sin \theta}} = \frac{r - ia \sin \theta}{\sqrt{\Sigma}} = \frac{\sqrt{\Sigma}}{r + ia \sin \theta}, \quad (\text{E.38})$$

$$Z = \dot{\theta} + \frac{i}{\Sigma} \left(\frac{L_z}{\sin \theta} - aE \sin \theta \right), \quad (\text{E.39})$$

and $K = K_{ab}k^ak^b$ is Carter's constant. It is quick to show that $Z\bar{Z} = \frac{K}{\Sigma^2}$, so $\frac{Z}{\bar{Z}} = \frac{\Sigma^2 Z^2}{K}$. We can therefore write (E.37) as

$$e^{i\beta} = \frac{e^{i\alpha}}{\sqrt{K}} \left[\Sigma\dot{\theta} + i \left(\frac{L_z}{\sin\theta} - aE \sin\theta \right) \right]. \quad (\text{E.40})$$

The transformation (E.31)–(E.33) can be achieved by performing a series of Lorentz transformations of the symmetrised Kinnersley tetrad $\{\hat{k}, \hat{n}, \hat{m}, \bar{\hat{m}}\}$. First, perform a boost in the (\hat{k}, \hat{n}) -plane and a spatial rotation in the $(\hat{m}, \bar{\hat{m}})$ -plane, such that $k_1 = e^{+\gamma}\hat{k}$, $n_1 = e^{-\gamma}\hat{n}$ and $m_1 = e^{+i\beta}\bar{\hat{m}}$. (Note that we have changed the handedness in the $(\hat{m}, \bar{\hat{m}})$ -plane.) Second, perform a null rotation which keeps n_1 fixed of the form $k_2 = k_1 + \bar{\mathcal{E}}m_1 + \mathcal{E}\bar{m}_1 + \mathcal{E}\bar{\mathcal{E}}n_1$, $n_2 = n_1$, $m_2 = m_1 + \mathcal{E}n_1$, with $\mathcal{E} = 1$. Third, perform a null rotation which keeps k_2 fixed of the form $k_3 = k_2$, $n_3 = n_2 + \bar{\mathcal{B}}m_1 + \mathcal{B}\bar{m}_1 + \mathcal{B}\bar{\mathcal{B}}k_1$, $m_3 = m_2 + \mathcal{B}k_2$, with $\mathcal{B} = -\frac{1}{2}$. Finally, perform a boost in the (k_3, n_3) -plane and a spatial rotation in the (m_3, \bar{m}_3) -plane, such that $\tilde{k} = Ak_3$, $\tilde{n} = \frac{1}{A}n_3$ and $\tilde{m} = -ie^{i\alpha}m_3$, where A and $e^{i\alpha}$ are defined in (E.35) and (E.38), respectively.

The relationships (E.31)–(E.33) permit us to calculate quantities in terms of the symmetrised Kinnersley tetrad $\{\hat{k}, \hat{n}, \hat{m}, \bar{\hat{m}}\}$, and map them to the “tilded” Marck tetrad $\{k, \tilde{n}, \tilde{m}, \bar{\tilde{m}}\}$. Transforming from the “tilded” tetrad to the parallel-transported tetrad using the Lorentz transformation (7.66) with null rotation parameter (7.67) is then straightforward.

As an example, let us calculate the Weyl scalars in the “tilded” frame by transforming from the Kinnersley tetrad. It is well known that Kerr spacetime is Petrov type D, i.e., it has two double principal null directions; these are precisely the legs \hat{k} and \hat{n} of the symmetrised Kinnersley tetrad. This means that, in the symmetrised Kinnersley tetrad, there is only one non-vanishing Weyl scalar:

$$\hat{\Psi}_2 = C_{abcd}\hat{k}^a\hat{m}^b\bar{\hat{m}}^c\hat{n}^d = -\frac{M}{(r - ia \cos\theta)^3}. \quad (\text{E.41})$$

One may now employ the sequence of Lorentz transformations described above to calculate the Weyl scalars in the “tilded” frame. We note that, in changing from the frame $\{\hat{k}, \hat{n}, \hat{m}, \bar{\hat{m}}\}$ to the frame $\{k_1, n_1, m_1, \bar{m}_1\}$ we interchanged m_1 and \bar{m}_1 ; the only non-vanishing Weyl scalar in the latter frame is therefore $\Psi_2^{(1)} = C_{abcd}k_1^am_1^b\bar{m}_1^cn_1^d = C_{abcd}\hat{k}^a\bar{\hat{m}}^b\hat{m}^c\hat{n}^d = \bar{\hat{\Psi}}_2$, which is the complex conjugate of (E.41). Applying the sequence of Lorentz transformations, we arrive at the Weyl scalars in the tilded basis:

$$\tilde{\Psi}_0 = C_{abcd}\tilde{k}^a\tilde{m}^b\tilde{k}^c\tilde{m}^d = -A^2e^{2i\alpha}(6\mathcal{E}^2)\bar{\hat{\Psi}}_2, \quad (\text{E.42})$$

$$\tilde{\Psi}_1 = C_{abcd}\tilde{k}^a\tilde{n}^b\tilde{k}^c\tilde{m}^d = -Aie^{i\alpha}(3\mathcal{E} + 6\mathcal{E}^2\bar{\mathcal{B}})\bar{\hat{\Psi}}_2, \quad (\text{E.43})$$

$$\tilde{\Psi}_2 = C_{abcd}\tilde{k}^a\tilde{m}^b\bar{\tilde{m}}^c\tilde{n}^d = (1 + 6\mathcal{E}\bar{\mathcal{B}} + 6\mathcal{E}^2\bar{\mathcal{B}}^2)\bar{\hat{\Psi}}_2, \quad (\text{E.44})$$

$$\tilde{\Psi}_3 = C_{abcd}\tilde{k}^a\tilde{n}^b\bar{\tilde{m}}^c\tilde{n}^d = \frac{ie^{-i\alpha}}{A}(3\bar{\mathcal{B}} + 9\mathcal{E}\bar{\mathcal{B}}^2 + 6\mathcal{E}^2\bar{\mathcal{B}}^3)\bar{\hat{\Psi}}_2, \quad (\text{E.45})$$

$$\tilde{\Psi}_4 = C_{abcd}\bar{\tilde{m}}^a\tilde{n}^b\bar{\tilde{m}}^c\tilde{n}^d = -\frac{e^{-2i\alpha}}{A^2}(6\bar{\mathcal{B}}^2 + 12\mathcal{E}\bar{\mathcal{B}}^3 + 6\mathcal{E}^2\bar{\mathcal{B}}^4)\bar{\hat{\Psi}}_2, \quad (\text{E.46})$$

where $\mathcal{E} = 1$ and $\mathcal{B} = -\frac{1}{2}$ are the null rotation parameters defined above. Upon simplification, the expressions agree (E.42)–(E.46) with (7.119)–(7.123), which were computed

by contracting the “tilded” tetrad legs with the Weyl tensor. Using the expressions for the Weyl scalars in the “tilded” frame, one may obtain the Weyl scalars in the parallel-transported basis using by applying the Lorentz transformation outlined in Section 7.5.2.

E.4 Calculating directional derivatives of Weyl scalars

In Section 7.5.3, we provide expressions for directional derivatives of the Weyl scalars Ψ_0 and Ψ_1 along the legs m^a and \bar{m}^a of our parallel-transported null tetrad. These involve directional derivatives of the Weyl tensor projected along tetrad legs, e.g. (7.131) features the quantity $(\delta C)_{kmkm} = (m^e \nabla_e C_{abcd}) k^a m^b k^c m^d$. In this section, we calculate these quantities explicitly, by first calculating analogous quantities in the Kinnersley tetrad, then transforming to the “tilded” complex null tetrad using the transformation (E.31)–(E.33) of Section E.3.

Let us denote directional derivatives along Kinnersley’s symmetrised tetrad (E.30) by

$$\hat{D} = \hat{k}^a \nabla_a, \quad \hat{\Delta} = \hat{n}^a \nabla_a, \quad \hat{\delta} = \hat{m}^a \nabla_a, \quad \bar{\hat{\delta}} = \bar{\hat{m}}^a \nabla_a. \quad (\text{E.47})$$

Using a symbolic algebra package, we find

$$(\hat{D}C_{abcd})k^a \tilde{m}^b k^c \tilde{m}^d = \frac{6iKM a \sin \theta}{\sqrt{2}\varrho^7} e^{-\gamma-i\beta} - \frac{9KM}{\varrho^6} \sqrt{\frac{\Delta}{2\Sigma}}, \quad (\text{E.48})$$

$$(\hat{\Delta}C_{abcd})k^a \tilde{m}^b k^c \tilde{m}^d = \frac{6iKM a \sin \theta}{\sqrt{2}\varrho^6 \bar{\varrho}} e^{+\gamma+i\beta} + \frac{9KM}{\varrho^6} \sqrt{\frac{\Delta}{2\Sigma}}, \quad (\text{E.49})$$

$$(\hat{\delta}C_{abcd})k^a \tilde{m}^b k^c \tilde{m}^d = \frac{9iKM a \sin \theta}{\sqrt{2}\varrho^7} - \frac{6KM e^{+\gamma+i\beta}}{\varrho^6} \sqrt{\frac{\Delta}{2\Sigma}}, \quad (\text{E.50})$$

$$(\bar{\hat{\delta}}C_{abcd})k^a \tilde{m}^b k^c \tilde{m}^d = \frac{9iKM a \sin \theta}{\sqrt{2}\varrho^6 \bar{\varrho}} + \frac{6KM e^{-\gamma-i\beta}}{\varrho^6} \sqrt{\frac{\Delta}{2\Sigma}}, \quad (\text{E.51})$$

$$(\hat{D}C_{abcd})k^a \tilde{m}^b \bar{m}^c \tilde{m}^d = -\frac{3AM a \sin \theta}{\sqrt{2}\varrho^5} e^{i\alpha-\gamma-i\beta}, \quad (\text{E.52})$$

$$(\hat{\Delta}C_{abcd})k^a \tilde{m}^b \bar{m}^c \tilde{m}^d = \frac{3AM a \sin \theta}{\sqrt{2}\varrho^4 \bar{\varrho}} e^{i\alpha+\gamma+i\beta}, \quad (\text{E.53})$$

$$(\hat{\delta}C_{abcd})k^a \tilde{m}^b \bar{m}^c \tilde{m}^d = \frac{3iAM}{\varrho^4} \sqrt{\frac{\Delta}{2\Sigma}} e^{i\alpha+\gamma+i\beta}, \quad (\text{E.54})$$

$$(\bar{\hat{\delta}}C_{abcd})k^a \tilde{m}^b \bar{m}^c \tilde{m}^d = \frac{3iAM}{\varrho^4} \sqrt{\frac{\Delta}{2\Sigma}} e^{i\alpha-\gamma-i\beta}, \quad (\text{E.55})$$

where we have introduced the quantity $\varrho = r + ia \cos \theta$ for notational simplicity. Trans-

forming to the “tilded” complex null tetrad via (E.31)–(E.33), we find

$$(DC_{abcd})k^a\tilde{m}^bk^c\tilde{m}^d = -\frac{15KM}{\varrho^6}(\dot{r} - i\dot{\theta}a\sin\theta) = D\Psi_0, \quad (\text{E.56})$$

$$(\tilde{\delta}C_{abcd})k^a\tilde{m}^bk^c\tilde{m}^d = \frac{15iM\sqrt{2K}}{2\varrho^7}[2a(L_z - aE) - E(r^2 - a^2\cos^2\theta)], \quad (\text{E.57})$$

$$(\bar{\delta}C_{abcd})k^a\tilde{m}^bk^c\tilde{m}^d = \frac{3iEM\sqrt{2K}}{2\varrho^5}, \quad (\text{E.58})$$

$$(DC_{abcd})k^a\tilde{m}^b\bar{m}^c\tilde{m}^d = \frac{3iEM\sqrt{2K}}{2\varrho^5}, \quad (\text{E.59})$$

$$(\tilde{\delta}C_{abcd})k^a\tilde{m}^b\bar{m}^c\tilde{m}^d = -\frac{3M\bar{\varrho}}{2\varrho^5}(\dot{r} + i\dot{\theta}a\sin\theta), \quad (\text{E.60})$$

$$(\bar{\delta}C_{abcd})k^a\tilde{m}^b\bar{m}^c\tilde{m}^d = -\frac{3M}{2\varrho^4}(\dot{r} - i\dot{\theta}a\sin\theta) = D\Psi_2. \quad (\text{E.61})$$

Finally, we may calculate the desired quantities by using a null rotation to transform to the parallel-transported complex null tetrad. Under this transformation, the leg \tilde{m}^a transforms as $m^a = \tilde{m}^a + Bk^a$, where the null rotation parameter B is defined in (7.67). Using this, we find

$$(\delta C)_{k m k m} = (\tilde{\delta} C)_{k \tilde{m} k \tilde{m}} + B(DC)_{k \tilde{m} k \tilde{m}}, \quad (\text{E.62})$$

$$(\bar{\delta} C)_{k m k m} = (\bar{\delta} C)_{k \tilde{m} k \tilde{m}} + B(DC)_{k \tilde{m} k \tilde{m}}, \quad (\text{E.63})$$

$$(\delta C)_{k m \bar{m} m} = (\tilde{\delta} C)_{k \tilde{m} \bar{m} \tilde{m}} + B(DC)_{k \tilde{m} \bar{m} \tilde{m}} + \bar{B}(\tilde{\delta} C)_{k \tilde{m} k \tilde{m}} + B\bar{B}(DC)_{k \tilde{m} k \tilde{m}}, \quad (\text{E.64})$$

$$(\bar{\delta} C)_{k m \bar{m} m} = (\bar{\delta} C)_{k \tilde{m} \bar{m} \tilde{m}} + \bar{B}(DC)_{k \tilde{m} \bar{m} \tilde{m}} + \bar{B}(\tilde{\delta} C)_{k \tilde{m} k \tilde{m}} + \bar{B}^2(DC)_{k \tilde{m} k \tilde{m}}, \quad (\text{E.65})$$

where explicit expressions for the quantities on the right-hand side are given in (E.56)–(E.61). The directional derivatives (E.62)–(E.65) can now be inserted into (7.131)–(7.134) to complete the derivation of $\delta\Psi_0$, $\bar{\delta}\Psi_0$, $\delta\Psi_1$ and $\bar{\delta}\Psi_1$.

Bibliography

- [1] The Event Horizon Telescope Collaboration, “First M87 Event Horizon Telescope results. I. The shadow of the supermassive black hole,” *The Astrophysical Journal Letters*, vol. 875, no. L1, p. 1, 2019.
- [2] The Event Horizon Telescope Collaboration, “First M87 Event Horizon Telescope Results. II. Array and Instrumentation,” *The Astrophysical Journal Letters*, vol. 875, no. L2, p. 1, 2019.
- [3] The Event Horizon Telescope Collaboration, “First M87 Event Horizon Telescope Results. III. Data Processing and Calibration,” *The Astrophysical Journal Letters*, vol. 875, no. L3, p. 1, 2019.
- [4] The Event Horizon Telescope Collaboration, “First M87 Event Horizon Telescope Results. IV. Imaging the Central Supermassive Black Hole,” *The Astrophysical Journal Letters*, vol. 875, no. L4, p. 1, 2019.
- [5] The Event Horizon Telescope Collaboration, “First M87 Event Horizon Telescope Results. V. Physical Origin of the Asymmetric Ring,” *The Astrophysical Journal Letters*, vol. 875, no. L5, p. 1, 2019.
- [6] The Event Horizon Telescope Collaboration, “First M87 Event Horizon Telescope Results. VI. The Shadow and Mass of the Central Black Hole,” *The Astrophysical Journal Letters*, vol. 875, no. L6, p. 1, 2019.
- [7] V. Perlick, “Gravitational Lensing from a Spacetime Perspective,” *Living Reviews in Relativity*, vol. 7, no. 1, p. 9, 2004.
- [8] D. Psaltis, R. Narayan, V. L. Fish, A. E. Broderick, A. Loeb, and S. S. Doeleman, “Event Horizon Telescope evidence for alignment of the black hole in the center of the Milky Way with the inner stellar disk,” *The Astrophysical Journal*, vol. 798, no. 1, p. 15, 2014.
- [9] T. Johannsen and D. Psaltis, “Testing the no-hair theorem with observations in the electromagnetic spectrum. II. Black hole images,” *The Astrophysical Journal*, vol. 718, no. 1, p. 446, 2010.
- [10] A. E. Broderick, T. Johannsen, A. Loeb, and D. Psaltis, “Testing the no-hair theorem with event horizon telescope observations of Sagittarius A*,” *The Astrophysical Journal*, vol. 784, no. 1, p. 7, 2014.

- [11] T. Johannsen, “Testing the no-hair theorem with observations of black holes in the electromagnetic spectrum,” *Classical and Quantum Gravity*, vol. 33, no. 12, p. 124001, 2016.
- [12] A. Ricarte and J. Dexter, “The Event Horizon Telescope: exploring strong gravity and accretion physics,” *Monthly Notices of the Royal Astronomical Society*, vol. 446, no. 2, p. 1973, 2014.
- [13] T. Johannsen, A. E. Broderick, P. M. Plewa, S. Chatzopoulos, S. S. Doeleman, F. Eisenhauer, V. L. Fish, R. Genzel, O. Gerhard, and M. D. Johnson, “Testing General Relativity with the Shadow Size of Sgr A*,” *Physical Review Letters*, vol. 116, no. 3, p. 031101, 2016.
- [14] B. P. Abbott *et al.*, “Observation of Gravitational Waves from a Binary Black Hole Merger,” *Physical Review Letters*, vol. 116, no. 6, p. 061102, 2016.
- [15] B. P. Abbott *et al.*, “GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral,” *Physical Review Letters*, vol. 119, no. 16, p. 161101, 2017.
- [16] B. P. Abbott *et al.*, “Multi-messenger Observations of a Binary Neutron Star Merger,” *The Astrophysical Journal Letters*, vol. 848, no. 2, p. L12, 2017.
- [17] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, *et al.*, “Binary black hole mergers in the first advanced LIGO observing run,” *Physical Review X*, vol. 6, no. 4, p. 041015, 2016.
- [18] B. P. Abbott *et al.*, “GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs,” arXiv preprint arXiv:1811.12907 [astro-ph.HE], 2018.
- [19] The LSC–Virgo Data Analysis Council, “The LSC–Virgo White Paper on Gravitational Wave Data Analysis and Astrophysics (Summer 2018 edition).” LIGO Document T1800058-v2, 2018.
- [20] Y. Aso, Y. Michimura, K. Somiya, M. Ando, O. Miyakawa, T. Sekiguchi, D. Tatum, H. Yamamoto, *et al.*, “Interferometer design of the KAGRA gravitational wave detector,” *Physical Review D*, vol. 88, no. 4, p. 043007, 2013.
- [21] P. Amaro-Seoane, H. Audley, S. Babak, J. Baker, E. Barausse, P. Bender, E. Berti, P. Binetruy, M. Born, D. Bortoluzzi, *et al.*, “Laser Interferometer Space Antenna,” arXiv preprint arXiv:1702.00786 [astro-ph.IM], 2017.
- [22] J. Hough, “Long baseline gravitational wave detectors—Status and developments,” *Progress in Particle and Nuclear Physics*, vol. 66, no. 2, p. 233, 2011.
- [23] D. Psaltis, “Testing general relativity with the Event Horizon Telescope,” arXiv preprint arXiv:1806.09740 [gr-qc], 2018. To appear in “Testing the Kerr spacetime

with gravitational-wave and electromagnetic observations” (E. Berti, ed.), *General Relativity and Gravitation*.

- [24] K. Schwarzschild, “Über das Gravitationsfeld eines Massenpunktes nach der Einsteinschen Theorie,” *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften*, p. 189, 1916.
- [25] H. Reissner, “Über die Eigengravitation des elektrischen Feldes nach der Einsteinschen Theorie,” *Annalen der Physik*, vol. 355, no. 9, p. 106, 1916.
- [26] G. Nordström, “On the energy of the gravitational field in Einstein’s theory,” in *Proceedings of the Koninklijke Nederlandse Akademie van Wetenschappen*, vol. 20, p. 1238, 1918.
- [27] G. D. Birkhoff and R. E. Langer, *Relativity and Modern Physics*. Harvard University Press, 1923.
- [28] J. T. Jebsen, “Über die allgemeinen kugelsymmetrischen Lösungen der Einsteinschen Gravitationsgleichungen im Vakuum,” *Arkiv för Matematik, Astronomi och Fysik*, vol. 15, no. 18, 1921.
- [29] C. W. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation*. Freeman, 1973.
- [30] L. D. Landau, “On the theory of stars,” *Physikalische Zeitschrift der Sowjetunion*, vol. 1, no. 285, p. 152, 1932.
- [31] W. Baade and F. Zwicky, “On super-novae,” *Proceedings of the National Academy of Sciences*, vol. 20, no. 5, p. 254, 1934.
- [32] F. Zwicky, “On collapsed neutron stars,” *The Astrophysical Journal*, vol. 88, p. 522, 1938.
- [33] J. R. Oppenheimer and G. M. Volkoff, “On massive neutron cores,” *Physical Review*, vol. 55, no. 4, p. 374, 1939.
- [34] J. R. Oppenheimer and H. Snyder, “On continued gravitational contraction,” *Physical Review*, vol. 56, no. 5, p. 455, 1939.
- [35] R. P. Kerr, “Gravitational field of a spinning mass as an example of algebraically special metrics,” *Physical review letters*, vol. 11, no. 5, p. 237, 1963.
- [36] E. T. Newman, E. Couch, K. Chinnapared, A. Exton, A. Prakash, and R. Torrence, “Metric of a rotating, charged mass,” *Journal of Mathematical Physics*, vol. 6, no. 6, p. 918, 1965.
- [37] E. T. Newman and A. I. Janis, “Note on the Kerr spinning-particle metric,” *Journal of Mathematical Physics*, vol. 6, no. 6, p. 915, 1965.
- [38] W. Israel, “Event horizons in static vacuum space-times,” *Physical Review*, vol. 164, no. 5, p. 1776, 1967.

- [39] W. Israel, “Event horizons in static electrovac space-times,” *Communications in Mathematical Physics*, vol. 8, no. 3, p. 245, 1968.
- [40] D. C. Robinson, “Uniqueness of the Kerr black hole,” *Physical Review Letters*, vol. 34, no. 14, p. 905, 1975.
- [41] R. Penrose, “Gravitational Collapse: The Role of General Relativity,” *General Relativity and Gravitation*, vol. 34, no. 7, p. 1141, 2002.
- [42] S. W. Hawking and R. Penrose, “The singularities of gravitational collapse and cosmology,” *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, vol. 314, no. 1519, p. 529, 1970.
- [43] J. M. Bardeen, B. Carter, and S. W. Hawking, “The four laws of black hole mechanics,” *Communications in Mathematical Physics*, vol. 31, no. 2, p. 161, 1973.
- [44] J. D. Bekenstein, “Black holes and entropy,” *Physical Review D*, vol. 7, no. 8, p. 2333, 1973.
- [45] S. W. Hawking, “Black hole explosions?,” *Nature*, vol. 248, no. 5443, p. 30, 1974.
- [46] R. H. Boyer and R. W. Lindquist, “Maximal analytic extension of the Kerr metric,” *Journal of Mathematical Physics*, vol. 8, no. 2, p. 265, 1967.
- [47] B. Carter, “Hamilton-Jacobi and Schrodinger separable solutions of Einstein’s equations,” *Communications in Mathematical Physics*, vol. 10, no. 4, p. 280, 1968.
- [48] R. Floyd, *The dynamics of Kerr fields*. PhD thesis, University of London, 1973.
- [49] R. Penrose, “Naked singularities,” *Annals of the New York Academy of Sciences*, vol. 224, no. 1, p. 125, 1973.
- [50] V. P. Frolov, P. Krtouš, and D. Kubizňák, “Black holes, hidden symmetries, and complete integrability,” *Living Reviews in Relativity*, vol. 20, no. 1, p. 6, 2017.
- [51] J.-A. Marck, “Solution to the equations of parallel transport in Kerr geometry; tidal tensor,” *Proceedings of the Royal Society of London, Series A*, vol. 385, no. 1789, p. 431, 1983.
- [52] J.-A. Marck, “Parallel-tetrad on null geodesics in Kerr-Newman space-time,” *Physics Letters A*, vol. 97, no. 4, p. 140, 1983.
- [53] D. R. Brill, P. L. Chrzanowski, C. M. Pereira, E. D. Fackerell, and J. R. Ipser, “Solution of the scalar wave equation in a Kerr background by separation of variables,” *Physical Review D*, vol. 5, no. 8, p. 1913, 1972.
- [54] W. Unruh, “Separability of the neutrino equations in a Kerr background,” *Physical Review Letters*, vol. 31, no. 20, p. 1265, 1973.
- [55] S. Chandrasekhar, “The solution of Dirac’s equation in Kerr geometry,” *Proceedings of the Royal Society of London, Series A*, vol. 349, no. 1659, p. 571, 1976.

- [56] R. Güven, “Black holes have no superhair,” *Physical Review D*, vol. 22, no. 10, p. 2327, 1980.
- [57] V. P. Frolov, P. Krtouš, D. Kubizňák, and J. E. Santos, “Massive vector fields in rotating black-hole spacetimes: separability and quasinormal modes,” *Physical Review Letters*, vol. 120, no. 23, p. 231103, 2018.
- [58] S. A. Teukolsky, “Rotating Black Holes: Separable Wave Equations for Gravitational and Electromagnetic Perturbations,” *Physical Review Letters*, vol. 29, no. 16, p. 1114, 1972.
- [59] S. A. Teukolsky, “Perturbations of a rotating black hole. I. Fundamental equations for gravitational, electromagnetic and neutrino field perturbations,” *The Astrophysical Journal*, vol. 185, p. 635, 1973.
- [60] P. V. P. Cunha and C. A. R. Herdeiro, “Shadows and strong gravitational lensing: a brief review,” *General Relativity and Gravitation*, vol. 50, no. 4, p. 42, 2018.
- [61] A. S. Eddington, “A Comparison of Whitehead’s and Einstein’s Formulæ,” *Nature*, vol. 113, no. 2832, p. 192, 1924.
- [62] L. C. B. Crispino and D. J. Kennefick, “A hundred years of the first experimental test of general relativity,” *Nature Physics*, vol. 15, no. 5, p. 416, 2019.
- [63] C. M. Will, “Henry Cavendish, Johann von Soldner, and the deflection of light,” *American Journal of Physics*, vol. 56, no. 5, p. 413, 1988.
- [64] O. Chwolson, “Über eine mögliche Form fiktiver Doppelsterne,” *Astronomische Nachrichten*, vol. 221, p. 329, 1924.
- [65] A. Einstein, “Lens-like action of a star by the deviation of light in the gravitational field,” *Science*, vol. 84, no. 2188, p. 506, 1936.
- [66] S. Refsdal, “The gravitational lens effect,” *Monthly Notices of the Royal Astronomical Society*, vol. 128, no. 4, p. 295, 1964.
- [67] Y. G. Klimov, “The deflection of light rays in the gravitational fields of galaxies,” *Soviet Physics Doklady*, vol. 8, p. 119, 1963.
- [68] S. Liebes Jr., “Gravitational Lenses,” *Physical Review*, vol. 133, no. 3B, p. B835, 1964.
- [69] D. Walsh, R. F. Carswell, and R. J. Weymann, “0957+ 561 A, B: twin quasistellar objects or gravitational lens?,” *Nature*, vol. 279, no. 5712, p. 381, 1979.
- [70] J. Wambsganss, “Gravitational Lensing in Astronomy,” *Living Reviews in Relativity*, vol. 1, no. 1, p. 12, 1998.
- [71] S. Dodelson, *Gravitational Lensing*. Cambridge University Press, 2017.

- [72] J. L. Synge, “The escape of photons from gravitationally intense stars,” *Monthly Notices of the Royal Astronomical Society*, vol. 131, no. 3, p. 463, 1966.
- [73] J.-P. Luminet, “Image of a spherical black hole with thin accretion disk,” *Astronomy and Astrophysics*, vol. 75, no. 1–2, p. 228, 1979.
- [74] J. M. Bardeen, “Timelike and null geodesics in the Kerr metric,” in *Black Holes* (C. DeWitt and B. S. DeWitt, eds.), p. 215, Gordon and Breach, 1973.
- [75] N. A. Sharp, “Geodesics in black hole space-times,” *General Relativity and Gravitation*, vol. 10, no. 8, p. 659, 1979.
- [76] E. Teo, “Spherical photon orbits around a Kerr black hole,” *General Relativity and Gravitation*, vol. 35, no. 11, p. 1909, 2003.
- [77] E. P. T. Liang, “Equatorial circular orbits of some static gravitational fields with naked singularities,” *Physical Review D*, vol. 9, no. 12, p. 3257, 1974.
- [78] V. Balek, J. Bičák, and Z. Stuchlík, “The motion of the charged particles in the field of rotating charged black holes and naked singularities. II. The motion in the equatorial plane,” *Bulletin of the Astronomical Institutes of Czechoslovakia*, vol. 40, p. 133, 1989.
- [79] P. V. P. Cunha, C. A. R. Herdeiro, and E. Radu, “Fundamental photon orbits: Black hole shadows and spacetime instabilities,” *Physical Review D*, vol. 96, no. 2, p. 024039, 2017.
- [80] H. Falcke, F. Melia, and E. Agol, “Viewing the shadow of the black hole at the galactic center,” *The Astrophysical Journal Letters*, vol. 528, no. 1, p. L13, 1999.
- [81] S. R. Dolan, “Higher-order geometrical optics for electromagnetic waves on a curved spacetime,” arXiv preprint arXiv:1801.02273 [gr-qc], 2018.
- [82] R. M. Wald, *General Relativity*. University of Chicago Press, 1984.
- [83] H. Stephani, D. Kramer, M. MacCallum, C. Hoenselaers, and E. Herlt, *Exact Solutions of Einstein’s Field Equations*. Cambridge Monographs on Mathematical Physics, Cambridge University Press, 2003.
- [84] V. I. Arnol’d, *Mathematical Methods of Classical Mechanics*, vol. 60 of *Graduate Texts in Mathematics*. Springer-Verlag, 1989.
- [85] F. J. Ernst, “New Formulation of the Axially Symmetric Gravitational Field Problem. II,” *Physical Review*, vol. 168, no. 5, p. 1415, 1968.
- [86] J. B. Griffiths and J. Podolský, *Exact Space-Times in Einstein’s General Relativity*. Cambridge Monographs on Mathematical Physics, Cambridge University Press, 2009.
- [87] D. L. Wiltshire, M. Visser, and S. M. Scott, eds., *The Kerr Spacetime: Rotating Black Holes in General Relativity*. Cambridge University Press, 2009.

- [88] R. P. Kerr, “Gravitational collapse and rotation,” in *Quasi-stellar sources and gravitational collapse: including the proceedings of the First Texas Symposium on Relativistic Astrophysics* (I. Robinson, A. Schild, and E. L. Schücking, eds.), p. 99, University of Chicago Press, 1965.
- [89] S. Chandrasekhar, “The two-centre problem in general relativity: the scattering of radiation by two extreme Reissner–Nordström black-holes,” *Proceedings of the Royal Society of London, Series A*, vol. 421, no. 1861, p. 227, 1989.
- [90] H. Weyl, “Zur Gravitationstheorie,” *Annalen der Physik*, vol. 359, no. 18, p. 117, 1917.
- [91] G. B. Jeffery, “The field of an electron on Einstein’s theory of gravitation,” *Proceedings of the Royal Society of London, Series A*, vol. 99, no. 697, p. 123, 1921.
- [92] M. Heusler, *Black Hole Uniqueness Theorems*. Cambridge University Press, 1996.
- [93] E. Newman and R. Penrose, “An approach to gravitational radiation by a method of spin coefficients,” *Journal of Mathematical Physics*, vol. 3, no. 3, p. 566, 1962.
- [94] R. Geroch, A. Held, and R. Penrose, “A space-time calculus based on pairs of null directions,” *Journal of Mathematical Physics*, vol. 14, no. 7, p. 874, 1973.
- [95] R. L. Devaney, *An Introduction to Chaotic Dynamical Systems*. Addison-Wesley, 1989.
- [96] P. Glendinning, *Stability, Instability and Chaos: an introduction to the theory of nonlinear differential equations*. Cambridge University Press, 1994.
- [97] K. T. Alligood, T. D. Sauer, and J. A. Yorke, *Chaos: An Introduction to Dynamical Systems*. Textbooks in Mathematical Sciences, Springer-Verlag, 1996.
- [98] M. Brin and G. Stuck, *Introduction to Dynamical Systems*. Cambridge University Press, 2002.
- [99] E. Ott, *Chaos in Dynamical Systems*. Cambridge University Press, 2002.
- [100] P. Gaspard, *Chaos, Scattering and Statistical Mechanics*. Cambridge University Press, 2005.
- [101] A. M. Ozorio de Almeida, *Hamiltonian Systems: Chaos and Quantization*. Cambridge University Press, 1990.
- [102] J. H. Lowenstein, *Essentials of Hamiltonian Dynamics*. Cambridge University Press, 2012.
- [103] M. V. Berry, “Regular and irregular motion,” in *AIP Conference Proceedings*, vol. 46, p. 16, 1978.
- [104] M. Cariglia, “Hidden symmetries of dynamics in classical and quantum physics,” *Reviews of Modern Physics*, vol. 86, no. 4, p. 1283, 2014.

- [105] P. Krtouš, “Electromagnetic field in higher-dimensional black-hole spacetimes,” *Physical Review D*, vol. 76, no. 8, p. 084035, 2007.
- [106] P. Krtouš, V. P. Frolov, and D. Kubizňák, “Hidden symmetries of higher-dimensional black holes and uniqueness of the Kerr-NUT-(A) dS spacetime,” *Physical Review D*, vol. 78, no. 6, p. 064022, 2008.
- [107] T. Houri, T. Oota, and Y. Yasui, “Closed conformal Killing–Yano tensor and the uniqueness of generalized Kerr–NUT–de Sitter spacetime,” *Classical and Quantum Gravity*, vol. 26, no. 4, p. 045015, 2009.
- [108] Y. Yasui and T. Houri, “Hidden symmetry and exact solutions in Einstein gravity,” *Progress of Theoretical Physics Supplement*, vol. 189, p. 126, 2011.
- [109] V. P. Frolov and D. Kubizňák, “Hidden symmetries of higher-dimensional rotating black holes,” *Physical Review Letters*, vol. 98, no. 1, p. 011101, 2007.
- [110] D. Kubizňák and V. P. Frolov, “The hidden symmetry of higher dimensional Kerr–NUT–AdS spacetimes,” *Classical and Quantum Gravity*, vol. 24, no. 3, p. F1, 2007.
- [111] T. Houri, T. Oota, and Y. Yasui, “Closed conformal Killing–Yano tensor and Kerr–NUT–de Sitter space–time uniqueness,” *Physics Letters B*, vol. 656, no. 4–5, p. 214, 2007.
- [112] J. Aguirre, R. L. Viana, and M. A. F. Sanjuán, “Fractal structures in nonlinear dynamics,” *Reviews of Modern Physics*, vol. 81, no. 1, p. 333, 2009.
- [113] K. Falconer, *Techniques in Fractal Geometry*. John Wiley & Sons, 1997.
- [114] K. Falconer, *Fractal Geometry: Mathematical Foundations and Applications*. John Wiley & Sons, 2004.
- [115] H. Poincaré, *Les Méthodes Nouvelles de la Mécanique Céleste*, vol. I–III. Gauthier-Villars, 1899.
- [116] M. Tabor, *Chaos and Integrability in Nonlinear Dynamics*. John Wiley & Sons, 1989.
- [117] M. Hénon and C. Heiles, “The Applicability of the Third Integral of Motion: Some Numerical Experiments,” *The Astronomical Journal*, vol. 69, p. 73, 1964.
- [118] E. Ott and T. Tél, “Chaotic scattering: An introduction,” *Chaos*, vol. 3, no. 4, p. 417, 1993.
- [119] B. Eckhardt, “Fractal properties of scattering singularities,” *Journal of Physics A*, vol. 20, no. 17, p. 5971, 1987.
- [120] P. Gaspard and S. A. Rice, “Scattering from a classically chaotic repeller,” *The Journal of Chemical Physics*, vol. 90, no. 4, p. 2225, 1989.
- [121] B. Eckhardt, “Irregular scattering,” *Physica D*, vol. 33, no. 1–3, p. 89, 1988.

- [122] J. M. Seoane and M. A. F. Sanjuán, “New developments in classical chaotic scattering,” *Reports on Progress in Physics*, vol. 76, no. 1, p. 016001, 2012.
- [123] E. Ott, “Strange attractors and chaotic motions of dynamical systems,” *Reviews of Modern Physics*, vol. 53, no. 4, p. 655, 1981.
- [124] J. Milnor, “On the Concept of Attractor,” in *The Theory of Chaotic Attractors* (B. R. Hunt, T. Y. Li, J. A. Kennedy, and H. E. Nusse, eds.), p. 243, Springer, 1985.
- [125] J.-P. Eckmann and D. Ruelle, “Ergodic theory of chaos and strange attractors,” in *The Theory of Chaotic Attractors* (B. R. Hunt, T. Y. Li, J. A. Kennedy, and H. E. Nusse, eds.), p. 273, Springer, 1985.
- [126] D. F. Chernoff and J. D. Barrow, “Chaos in the mixmaster universe,” *Physical Review Letters*, vol. 50, no. 2, p. 134, 1983.
- [127] D. Hobill, A. Burd, and A. A. Coley, eds., *Deterministic Chaos in General Relativity*, vol. 332 of *Nato Science Series B*. Springer, 1994.
- [128] N. J. Cornish and J. J. Levin, “The mixmaster universe is chaotic,” *Physical Review Letters*, vol. 78, no. 6, p. 998, 1997.
- [129] O. Semerák and P. Suková, “Free motion around black holes with discs or rings: between integrability and chaos – I,” *Monthly Notices of the Royal Astronomical Society*, vol. 404, no. 2, p. 545, 2010.
- [130] S. Suzuki and K.-i. Maeda, “Chaos in Schwarzschild spacetime: The motion of a spinning particle,” *Physical Review D*, vol. 55, no. 8, p. 4848, 1997.
- [131] M. D. Hartl, “Dynamics of spinning test particles in Kerr spacetime,” *Physical Review D*, vol. 67, no. 2, p. 024005, 2003.
- [132] M. D. Hartl, “Survey of spinning test particle orbits in Kerr spacetime,” *Physical Review D*, vol. 67, no. 10, p. 104023, 2003.
- [133] L. Bombelli and E. Calzetta, “Chaos around a black hole,” *Classical and Quantum Gravity*, vol. 9, no. 12, p. 2573, 1992.
- [134] J. Levin, “Gravity waves, chaos, and spinning compact binaries,” *Physical Review Letters*, vol. 84, no. 16, p. 3515, 2000.
- [135] P. S. Letelier and W. M. Vieira, “Chaos in black holes surrounded by gravitational waves,” *Classical and Quantum Gravity*, vol. 14, no. 5, p. 1249, 1997.
- [136] N. J. Cornish, “Chaos and gravitational waves,” *Physical Review D*, vol. 64, no. 8, p. 084011, 2001.
- [137] W. M. Vieira and P. S. Letelier, “Chaos around a Hénon-Heiles-inspired exact perturbation of a black hole,” *Physical Review Letters*, vol. 76, no. 9, p. 1409, 1996.

- [138] A. P. S. de Moura and P. S. Letelier, “Scattering map for two black holes,” *Physical Review E*, vol. 62, no. 4, p. 4784, 2000.
- [139] Y. Sota, S. Suzuki, and K.-i. Maeda, “Chaos in static axisymmetric spacetimes: I. Vacuum case,” *Classical and Quantum Gravity*, vol. 13, no. 5, p. 1241, 1996.
- [140] G. Contopoulos, “Periodic orbits and chaos around two black holes,” *Proceedings of the Royal Society of London, Series A*, vol. 431, no. 1881, p. 183, 1990.
- [141] G. Contopoulos, “Periodic orbits and chaos around two fixed black holes. II,” *Proceedings of the Royal Society of London, Series A*, vol. 435, no. 1895, p. 551, 1991.
- [142] G. Contopoulos, *Order and Chaos in Dynamical Astronomy*. Springer-Verlag Berlin Heidelberg, 2002.
- [143] I. Newton, *Opticks. A Treatise of the Reflections, Refractions, Inflections & Colours of Light*. Dover, 1979. Republication based on the 4th edition, 1730.
- [144] J. Michell, “On the means of discovering the distance, magnitude, &c. of the fixed stars, in consequence of the diminution of the velocity of their light, in case such a diminution should be found to take place in any of them, and such other data should be procured from observations, as would be farther necessary for that purpose. By the Rev. John Michell, BDFRS In a letter to Henry Cavendish, Esq. FRS and AS,” *Philosophical Transactions of the Royal Society of London*, no. 74, p. 35, 1784.
- [145] H. Cavendish, *The Scientific Papers of the Honourable Henry Cavendish, F.R.S.*, vol. 2. Cambridge University Press, 2011.
- [146] S. Schaffer, “John Michell and Black Holes,” *Journal for the History of Astronomy*, vol. 10, p. 42, 1979.
- [147] P.-S. Laplace, *Exposition du système du monde*. Cambridge University Press, 2009. Republication based on the 2nd edition, 1799.
- [148] J. G. von Soldner, “Ueber die Ablenkung eines Lichtstrals von seiner geradlinigen Bewegung, durch die Attraktion eines Weltkörpers, an welchem er nahe vorbei geht,” *Astronomisches Jahrbuch für das Jahr 1804*, p. 161, 1801.
- [149] S. L. Jaki, “Johann Georg von Soldner and the gravitational bending of light, with an English translation of his essay on it published in 1801,” *Foundations of Physics*, vol. 8, no. 11–12, p. 927, 1978.
- [150] A. Einstein, “Über den Einfluß der Schwerkraft auf die Ausbreitung des Lichtes,” *Annalen der Physik*, vol. 14, p. 898, 1911.
- [151] A. Einstein, “Erklärung der Perihelionbewegung der Merkur aus der allgemeinen Relativitätstheorie,” *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften*, vol. 47, no. 2, p. 831, 1915.

- [152] C. M. Will, “The Confrontation Between General Relativity and Experiment,” *Living Reviews in Relativity*, vol. 9, no. 1, p. 3, 2006.
- [153] C. G. Darwin, “The gravity field of a particle,” *Proceedings of the Royal Society of London, Series A*, vol. 249, no. 1257, p. 180, 1959.
- [154] G. S. Bisnovatyi-Kogan and O. Y. Tsupko, “Strong gravitational lensing by Schwarzschild black holes,” *Astrophysics*, vol. 51, no. 1, p. 99, 2008.
- [155] F. W. Dyson, A. S. Eddington, and C. Davidson, “A determination of the deflection of light by the Sun’s gravitational field, from observations made at the total eclipse of May 29, 1919,” *Philosophical Transactions of the Royal Society of London, Series A*, vol. 220, no. 571–581, p. 291, 1920.
- [156] S. S. Shapiro, J. L. Davis, D. E. Lebach, and J. S. Gregory, “Measurement of the solar gravitational deflection of radio waves using geodetic very-long-baseline interferometry data, 1979–1999,” *Physical Review Letters*, vol. 92, no. 12, p. 121101, 2004.
- [157] D. G. Bruns, “Gravitational starlight deflection measurements during the 21 August 2017 total solar eclipse,” *Classical and Quantum Gravity*, vol. 35, no. 7, p. 075009, 2018.
- [158] F. Zwicky, “Nebulae as Gravitational Lenses,” *Physical Review*, vol. 51, no. 4, p. 290, 1937.
- [159] S. Refsdal, “On the possibility of determining Hubble’s parameter and the masses of galaxies from the gravitational lens effect,” *Monthly Notices of the Royal Astronomical Society*, vol. 128, no. 4, p. 307, 1964.
- [160] R. Sachs, “Gravitational waves in general relativity. VI. The outgoing radiation condition,” *Proceedings of the Royal Society of London, Series A*, vol. 264, no. 1318, p. 309, 1961.
- [161] M. Schmidt, “3C 273: A Star-Like Object with Large Red-Shift,” *Nature*, vol. 197, p. 1040, 1963.
- [162] P. Schneider, J. Ehlers, and E. E. Falco, *Gravitational Lenses*. Springer-Verlag, 1992.
- [163] P. Schneider, C. S. Kochanek, and J. Wambsganss, *Gravitational Lensing: Strong, Weak and Micro*. Saas-Fee Advanced Course 33, Springer-Verlag, 2006.
- [164] R. Narayan and M. Bartelmann, “Lectures on gravitational lensing,” in *Proceedings of the 1995 Jerusalem Winter School in Theoretical Physics: Formation of Structure in the Universe* (A. Dekel and J. P. Ostriker, eds.), 1996.
- [165] S. M. Carroll, *Spacetime and Geometry*. Cambridge University Press, 2019.

- [166] N. Inada, M. Oguri, B. Pindor, J. F. Hennawi, K. Chiu, W. Zheng, S.-I. Ichikawa, M. D. Gregg, R. H. Becker, Y. Suto, M. A. Strauss, E. L. Turner, C. R. Keeton, J. Annis, F. J. Castander, D. J. Eisenstein, J. A. Frieman, M. Fukugita, J. E. Gunn, D. E. Johnston, S. M. Kent, R. C. Nichol, G. T. Richards, H.-W. Rix, E. S. Sheldon, N. A. Bahcall, J. Brinkmann, Ž. Ivezić, D. Q. Lamb, T. A. McKay, D. P. Schneider, and D. G. York, “A gravitationally lensed quasar with quadruple images separated by 14.62 arcseconds,” *Nature*, vol. 426, no. 6968, p. 810, 2003.
- [167] E. Poisson and C. M. Will, *Gravity: Newtonian, Post-Newtonian, Relativistic*. Cambridge University Press, 2014.
- [168] J. Kristian and R. K. Sachs, “Observations in cosmology,” *The Astrophysical Journal*, vol. 143, p. 379, 1966. Republication: *General Relativity and Gravitation* vol. 43, no. 1, p. 337, 2011.
- [169] J. Ehlers, “Zum Übergang von der Wellenoptik zur geometrischen Optik in der allgemeinen Relativitätstheorie,” *Zeitschrift für Naturforschung A*, vol. 22, no. 9, p. 1328, 1967.
- [170] A. M. Anile, “Geometrical optics in general relativity: A study of the higher order corrections,” *Journal of Mathematical Physics*, vol. 17, no. 4, p. 576, 1976.
- [171] S. R. Dolan, “Geometrical optics for scalar, electromagnetic and gravitational waves on curved spacetime,” *International Journal of Modern Physics D*, vol. 27, no. 11, p. 1843010, 2017.
- [172] E. Poisson, *A Relativist’s Toolkit: The Mathematics of Black-Hole Mechanics*. Cambridge University Press, 2004.
- [173] P. Jordan, J. Ehlers, and R. K. Sachs, “Republication of: Contributions to the theory of pure gravitational radiation. Exact solutions of the field equations of the general theory of relativity II,” *General Relativity and Gravitation*, vol. 45, no. 12, p. 2691, 2013.
- [174] R. Kantowski, “Another interpretation of the optical scalars,” *Journal of Mathematical Physics*, vol. 9, no. 2, p. 336, 1968.
- [175] V. P. Frolov and I. D. Novikov, *Black Hole Physics: Basic Concepts and New Developments*, vol. 96 of *Fundamental Theories of Physics*. Kluwer Academic Publishers, 1998.
- [176] S. Chandrasekhar, *The Mathematical Theory of Black Holes*, vol. 69 of *International Series of Monographs on Physics*. Oxford University Press, 1983.
- [177] S. Hod, “Spherical null geodesics of rotating Kerr black holes,” *Physics Letters B*, vol. 718, no. 4–5, p. 1552, 2013.
- [178] C. T. Cunningham and J. M. Bardeen, “The optical appearance of a star orbiting an extreme Kerr black hole,” *The Astrophysical Journal*, vol. 183, p. 237, 1973.

- [179] C. T. Cunningham, “The effects of redshifts and focusing on the spectrum of an accretion disk around a Kerr black Hole,” *The Astrophysical Journal*, vol. 202, p. 788, 1975.
- [180] J. Dexter and E. Agol, “A fast new public code for computing photon orbits in a Kerr spacetime,” *The Astrophysical Journal*, vol. 696, no. 2, p. 1616, 2009.
- [181] T. Müller and F. Grave, “GeodesicViewer – A tool for exploring geodesics in the theory of relativity,” *Computer Physics Communications*, vol. 181, no. 2, p. 413, 2010.
- [182] F. H. Vincent, T. Paumard, E. Gourgoulhon, and G. Perrin, “GYOTO: a new general relativistic ray-tracing code,” *Classical and Quantum Gravity*, vol. 28, no. 22, p. 225011, 2011.
- [183] T. Müller, “GeoViS—Relativistic ray tracing in four-dimensional spacetimes,” *Computer Physics Communications*, vol. 185, no. 8, p. 2301, 2014.
- [184] P. V. P. Cunha, J. Grover, C. Herdeiro, E. Radu, H. Runarsson, and A. Wittig, “Chaotic lensing around boson stars and Kerr black holes with scalar hair,” *Physical Review D*, vol. 94, no. 10, p. 104023, 2016.
- [185] J.-A. Marck, “Short-cut method of solution of geodesic equations for Schwarzschild black hole,” *Classical and Quantum Gravity*, vol. 13, no. 3, p. 393, 1996.
- [186] J. Levin and G. Perez-Giz, “A periodic table for black hole orbits,” *Physical Review D*, vol. 77, no. 10, p. 103005, 2008.
- [187] T. Müller, “Analytic observation of a star orbiting a Schwarzschild black hole,” *General Relativity and Gravitation*, vol. 41, no. 3, p. 541, 2009.
- [188] S. C. Noble, P. K. Leung, C. F. Gammie, and L. G. Book, “Simulating the emission and outflows from accretion discs,” *Classical and Quantum Gravity*, vol. 24, no. 12, p. S259, 2007.
- [189] R. V. Shcherbakov and L. Huang, “General relativistic polarized radiative transfer: building a dynamics–observations interface,” *Monthly Notices of the Royal Astronomical Society*, vol. 410, no. 2, p. 1052, 2010.
- [190] A. Bohn, W. Thrope, F. Hébert, K. Henriksson, D. Bunandar, M. A. Scheel, and N. W. Taylor, “What does a binary black hole merger look like?,” *Classical and Quantum Gravity*, vol. 32, no. 6, p. 065002, 2015.
- [191] P. V. P. Cunha, C. A. R. Herdeiro, E. Radu, and H. F. Rúnarsson, “Shadows of Kerr black holes with scalar hair,” *Physical Review Letters*, vol. 115, no. 21, p. 211102, 2015.
- [192] B. P. Abbott *et al.*, “GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence,” *Physical Review Letters*, vol. 116, no. 24, p. 241103, 2016.

- [193] B. P. Abbott *et al.*, “GW170104: Observation of a 50-Solar-Mass Binary Black Hole Coalescence at Redshift 0.2,” *Physical Review Letters*, vol. 118, no. 22, p. 221101, 2017. Erratum: *Physical Review Letters*, vol. 121, no. 12, p. 129901, 2018.
- [194] B. P. Abbott *et al.*, “GW170608: Observation of a 19-solar-mass Binary Black Hole Coalescence,” *The Astrophysical Journal Letters*, vol. 851, no. 2, p. L35, 2017.
- [195] B. P. Abbott *et al.*, “GW170814: A Three-Detector Observation of Gravitational Waves from a Binary Black Hole Coalescence,” *Physical Review Letters*, vol. 119, no. 14, p. 141101, 2017.
- [196] S. Doeleman, E. Agol, D. Backer, F. Baganoff, G. C. Bower, A. Broderick, A. Fabian, V. Fish, C. Gammie, P. Ho, *et al.*, “Imaging an event horizon: submm-VLBI of a super massive black hole,” arXiv preprint arXiv:0906.3899 [gr-qc], 2009.
- [197] H. Falcke, “Imaging black holes: past, present and future,” in *Journal of Physics: Conference Series*, vol. 942, p. 012001, 2017.
- [198] D. Nitta, T. Chiba, and N. Sugiyama, “Shadows of colliding black holes,” *Physical Review D*, vol. 84, no. 6, p. 063008, 2011.
- [199] A. Yumoto, D. Nitta, T. Chiba, and N. Sugiyama, “Shadows of multi-black holes: analytic exploration,” *Physical Review D*, vol. 86, no. 10, p. 103001, 2012.
- [200] D. Kastor and J. Traschen, “Cosmological multi-black-hole solutions,” *Physical Review D*, vol. 47, no. 12, p. 5370, 1993.
- [201] J. O. Shipley and S. R. Dolan, “Binary black hole shadows, chaotic scattering and the Cantor set,” *Classical and Quantum Gravity*, vol. 33, no. 17, p. 175001, 2016.
- [202] S. D. Majumdar, “A Class of Exact Solutions of Einstein’s Field Equations,” *Physical Review*, vol. 72, p. 390, 1947.
- [203] A. Papapetrou, “A static solution of the equations of the gravitational field for an arbitrary charge-distribution,” *Proceedings of the Royal Irish Academy, Section A*, vol. 51, p. 191, 1947.
- [204] J. B. Hartle and S. W. Hawking, “Solutions of the Einstein-Maxwell equations with many black holes,” *Communications in Mathematical Physics*, vol. 26, no. 2, p. 87, 1972.
- [205] G. Contopoulos and H. Papadaki, “Newtonian and relativistic periodic orbits around two fixed black holes,” *Celestial Mechanics and Dynamical Astronomy*, vol. 55, no. 1, p. 47, 1993.
- [206] G. Contopoulos, N. Voglis, and C. Efthymiopoulos, “Chaos in relativity and cosmology,” in *International Astronomical Union Colloquium*, vol. 172, p. 1, Cambridge University Press, 1999.

- [207] G. Contopoulos and M. Harsoula, “Chaos in the case of two fixed black holes,” *Journal of Mathematical Physics*, vol. 45, no. 12, p. 4932, 2004.
- [208] G. Contopoulos and M. Harsoula, “Systems with escapes,” *Annals of the New York Academy of Sciences*, vol. 1045, no. 1, p. 139, 2005.
- [209] U. Yurtsever, “Geometry of chaos in the two-center problem in general relativity,” *Physical Review D*, vol. 52, no. 6, p. 3176, 1995.
- [210] C. P. Dettmann, N. E. Frankel, and N. J. Cornish, “Fractal basins and chaotic trajectories in multi-black-hole spacetimes,” *Physical Review D*, vol. 50, no. 2, p. R618, 1994.
- [211] C. P. Dettmann, N. E. Frankel, and N. J. Cornish, “Chaos and fractals around black holes,” *Fractals*, vol. 03, no. 01, p. 161, 1995.
- [212] N. J. Cornish and N. E. Frankel, “The black hole and the pea,” *Physical Review D*, vol. 56, no. 4, p. 1903, 1997.
- [213] N. J. Cornish and G. W. Gibbons, “A tale of two centres,” *Classical and Quantum Gravity*, vol. 14, no. 7, p. 1865, 1997.
- [214] J. Levin, “Chaos may make black holes bright,” *Physical Review D*, vol. 60, no. 6, p. 064015, 1999.
- [215] J. Levin, R. O’Reilly, and E. J. Copeland, “Gravity waves from homoclinic orbits of compact binaries,” *Physical Review D*, vol. 62, no. 2, p. 024023, 2000.
- [216] M. Szydłowski and A. Krawiec, “Description of chaos in simple relativistic systems,” *Physical Review D*, vol. 53, no. 12, p. 6893, 1996.
- [217] W. Hanan and E. Radu, “Chaotic motion in multi-black hole spacetimes and holographic screens,” *Modern Physics Letters A*, vol. 22, no. 06, p. 399, 2007.
- [218] D. Alonso, A. Ruiz, and M. Sánchez-Hernández, “Escape of photons from two fixed extreme Reissner-Nordström black holes,” *Physical Review D*, vol. 78, p. 104024, 2008.
- [219] C. D. Murray and S. F. Dermott, *Solar System Dynamics*. Cambridge University Press, 1999.
- [220] C. A. Coulson and A. Joseph, “A constant of the motion for the two-centre Kepler problem,” *International Journal of Quantum Chemistry*, vol. 1, no. 4, p. 337, 1967.
- [221] D. Lynden-Bell, “A simple derivation and interpretation of the third integral in stellar dynamics,” *Monthly Notices of the Royal Astronomical Society*, vol. 338, no. 1, p. 208, 2003.

- [222] R. Bach and H. Weyl, “Republication of: New solutions to Einstein’s equations of gravitation. B. Explicit determination of static, axially symmetric fields. By Rudolf Bach. With a supplement on the static two-body problem. By H. Weyl,” *General Relativity and Gravitation*, vol. 44, no. 3, p. 817, 2012.
- [223] F. S. Coelho and C. A. R. Herdeiro, “Relativistic Euler’s 3-body problem, optical geometry and the golden ratio,” *Physical Review D*, vol. 80, p. 104036, 2009.
- [224] C. M. Will, “Carter-like constants of the motion in Newtonian gravity and electrodynamics,” *Physical Review Letters*, vol. 102, p. 061101, 2009.
- [225] T. Assumpção, V. Cardoso, A. Ishibashi, M. Richartz, and M. Zilhao, “Black hole binaries: ergoregions, photon surfaces, wave scattering, and quasinormal modes,” *Physical Review D*, vol. 98, no. 6, p. 064036, 2018.
- [226] L. Poon, J. Campos, E. Ott, and C. Grebogi, “Wada basin boundaries in chaotic scattering,” *International Journal of Bifurcation and Chaos*, vol. 6, no. 2, p. 251, 1996.
- [227] J. Aguirre and M. A. F. Sanjuán, “Limit of small exits in open Hamiltonian systems,” *Physical Review E*, vol. 67, no. 5, p. 056201, 2003.
- [228] D. Finkelstein, “Past-future asymmetry of the gravitational field of a point particle,” *Physical Review*, vol. 110, no. 4, p. 965, 1958.
- [229] R. Penrose, “Republication of: Conformal treatment of infinity,” *General Relativity and Gravitation*, vol. 43, no. 3, p. 901, 2011.
- [230] B. Carter, “Complete analytic extension of the symmetry axis of Kerr’s solution of Einstein’s equations,” *Physical Review*, vol. 141, no. 4, p. 1242, 1966.
- [231] W. Israel and G. A. Wilson, “A class of stationary electromagnetic vacuum fields,” *Journal of Mathematical Physics*, vol. 13, no. 6, p. 865, 1972.
- [232] N. Bretón, V. S. Manko, and J. Aguilar-Sánchez, “On the equilibrium of charged masses in general relativity: the electrostatic case,” *Classical and Quantum Gravity*, vol. 15, no. 10, p. 3071, 1998.
- [233] V. S. Manko, J. Martín, and E. Ruiz, “Metric of two arbitrary Kerr–Newman sources located on the symmetry axis,” *Journal of Mathematical Physics*, vol. 35, no. 12, p. 6644, 1994.
- [234] P. V. P. Cunha, C. A. R. Herdeiro, and M. J. Rodríguez, “Shadows of exact binary black holes,” *Physical Review D*, vol. 98, no. 4, p. 044053, 2018.
- [235] D. Kramer and G. Neugebauer, “The superposition of two Kerr solutions,” *Physics Letters A*, vol. 75, no. 4, p. 259, 1980.

- [236] P. V. P. Cunha, C. A. R. Herdeiro, and M. J. Rodriguez, “Does the black hole shadow probe the event horizon geometry?,” *Physical Review D*, vol. 97, no. 8, p. 084020, 2018.
- [237] M. Wang, S. Chen, and J. Jing, “Shadows of Bonnor black dihole by chaotic lensing,” *Physical Review D*, vol. 97, no. 6, p. 064029, 2018.
- [238] S. Abdolrahimi, R. B. Mann, and C. Tzounis, “Distorted local shadows,” *Physical Review D*, vol. 91, no. 8, p. 084052, 2015.
- [239] M. Wang, S. Chen, and J. Jing, “Chaotic shadow of a non-Kerr rotating compact object with quadrupole mass moment,” *Physical Review D*, vol. 98, no. 10, p. 104040, 2018.
- [240] M. Wang, S. Chen, and J. Jing, “Effect of gravitational wave on shadow of a Schwarzschild black hole,” arXiv preprint arXiv:1908.04527 [gr-qc], 2019.
- [241] P. V. P. Cunha, C. A. R. Herdeiro, E. Radu, and H. F. Runarsson, “Shadows of Kerr black holes with and without scalar hair,” *International Journal of Modern Physics D*, vol. 25, no. 09, p. 1641021, 2016.
- [242] P. V. P. Cunha, J. A. Font, C. Herdeiro, E. Radu, N. Sanchis-Gual, and M. Zilhão, “Lensing and dynamics of ultracompact bosonic stars,” *Physical Review D*, vol. 96, no. 10, p. 104040, 2017.
- [243] L. Bernard, V. Cardoso, T. Ikeda, and M. Zilhão, “Physics of black hole binaries: Geodesics, relaxation modes, and energy extraction,” *Physical Review D*, vol. 100, p. 044002, 2019.
- [244] K. Yoneyama, “Theory of Continuous Set of Points,” *Tohoku Mathematical Journal*, vol. 12, p. 43, 1917.
- [245] J. Kennedy and J. A. Yorke, “Basins of Wada,” *Physica D*, vol. 51, no. 1–3, p. 213, 1991.
- [246] J. Aguirre, J. C. Vallejo, and M. A. F. Sanjuán, “Wada basins and chaotic invariant sets in the Hénon-Heiles system,” *Physical Review E*, vol. 64, no. 6, p. 066208, 2001.
- [247] J. Aguirre and M. A. F. Sanjuán, “Unpredictable behavior in the Duffing oscillator: Wada basins,” *Physica D*, vol. 171, no. 1–2, p. 41, 2002.
- [248] A. Daza, A. Wagemakers, and M. A. F. Sanjuán, “Ascertaining when a basin is Wada: the merging method,” *Scientific Reports*, vol. 8, no. 1, p. 9954, 2018.
- [249] A. Daza, J. O. Shipley, S. R. Dolan, and M. A. F. Sanjuán, “Wada structures in a binary black hole system,” *Physical Review D*, vol. 98, no. 8, p. 084050, 2018.
- [250] B. B. Mandelbrot, *The Fractal Geometry of Nature*. W. H. Freeman and Company, 1982. Revised edition of *Fractals*, 1977.

- [251] C. Grebogi, S. W. McDonald, E. Ott, and J. A. Yorke, “Final state sensitivity: an obstruction to predictability,” *Physics Letters A*, vol. 99, no. 9, p. 415, 1983.
- [252] S. W. McDonald, C. Grebogi, E. Ott, and J. A. Yorke, “Fractal basin boundaries,” *Physica D*, vol. 17, no. 2, p. 125, 1985.
- [253] J. G. Hocking and G. S. Young, *Topology*. Dover Publications, 1988.
- [254] L. E. J. Brouwer, “Zur Analysis Situs,” *Mathematische Annalen*, vol. 68, no. 3, p. 422, 1910.
- [255] H. E. Nusse and J. A. Yorke, “Wada basin boundaries and basin cells,” *Physica D*, vol. 90, no. 3, p. 242, 1996.
- [256] Z. Toroczkai, G. Károlyi, A. Péntek, T. Tél, C. Grebogi, and J. A. Yorke, “Wada dye boundaries in open hydrodynamical flows,” *Physica A*, vol. 239, no. 1–3, p. 235, 1997.
- [257] A. Daza, A. Wagemakers, M. A. F. Sanjuán, and J. A. Yorke, “Testing for basins of Wada,” *Scientific Reports*, vol. 5, p. 16579, 2015.
- [258] J. C. Alexander, J. A. Yorke, Z. You, and I. Kan, “Riddled basins,” *International Journal of Bifurcation and Chaos*, vol. 2, no. 04, p. 795, 1992.
- [259] V. Cardoso, E. Franzin, and P. Pani, “Is the gravitational-wave ringdown a probe of the event horizon?,” *Physical Review Letters*, vol. 116, no. 17, p. 171101, 2016. Erratum: *Physical Review Letters*, vol. 117, no. 8, p. 089902, 2016.
- [260] C. J. Goebel, “Comments on the “vibrations” of a Black Hole,” *The Astrophysical Journal*, vol. 172, p. L95, 1972.
- [261] B. Mashhoon, “Stability of charged rotating black holes in the eikonal approximation,” *Physical Review D*, vol. 31, no. 2, p. 290, 1985. Erratum: *Physical Review D*, vol. 31, no. 10, p. 2697, 1985.
- [262] V. Cardoso, A. S. Miranda, E. Berti, H. Witek, and V. T. Zanchin, “Geodesic stability, Lyapunov exponents and quasinormal modes,” *Physical Review D*, vol. 79, p. 064016, 2009.
- [263] S. R. Dolan, “Quasinormal mode spectrum of a Kerr black hole in the eikonal limit,” *Physical Review D*, vol. 82, p. 104003, 2010.
- [264] H. Yang, D. A. Nichols, F. Zhang, A. Zimmerman, Z. Zhang, and Y. Chen, “Quasinormal-mode spectrum of Kerr black holes and its geometric interpretation,” *Physical Review D*, vol. 86, p. 104006, 2012.
- [265] L. C. B. Crispino, S. R. Dolan, and E. S. Oliveira, “Electromagnetic Wave Scattering by Schwarzschild Black Holes,” *Physical Review Letters*, vol. 102, p. 231103, 2009.
- [266] R. A. Konoplya and A. Zhidenko, “Quasinormal modes of black holes: From astrophysics to string theory,” *Reviews of Modern Physics*, vol. 83, p. 793, 2011.

- [267] V. Cardoso, L. C. B. Crispino, C. F. B. Macedo, H. Okawa, and P. Pani, “Light rings as observational evidence for event horizons: long-lived modes, ergoregions and nonlinear instabilities of ultracompact objects,” *Physical Review D*, vol. 90, no. 4, p. 044069, 2014.
- [268] P. V. P. Cunha, E. Berti, and C. A. R. Herdeiro, “Light-Ring Stability for Ultracompact Objects,” *Physical Review Letters*, vol. 119, no. 25, p. 251102, 2017.
- [269] S. L. Liebling and C. Palenzuela, “Dynamical boson stars,” *Living Reviews in Relativity*, vol. 20, no. 1, p. 5, 2017.
- [270] R. Brito, V. Cardoso, C. A. R. Herdeiro, and E. Radu, “Proca stars: gravitating Bose–Einstein condensates of massive spin 1 particles,” *Physics Letters B*, vol. 752, p. 291, 2016.
- [271] P. O. Mazur and E. Mottola, “Gravitational vacuum condensate stars,” *Proceedings of the National Academy of Sciences*, vol. 101, no. 26, p. 9545, 2004.
- [272] C. Chirenti and L. Rezzolla, “Did GW150914 produce a rotating gravastar?,” *Physical Review D*, vol. 94, no. 8, p. 084016, 2016.
- [273] T. Damour and S. N. Solodukhin, “Wormholes as black hole foils,” *Physical Review D*, vol. 76, no. 2, p. 024016, 2007.
- [274] J. L. Friedman, “Ergosphere instability,” *Communications in Mathematical Physics*, vol. 63, no. 3, p. 243, 1978.
- [275] P. Pani, E. Barausse, E. Berti, and V. Cardoso, “Gravitational instabilities of super-spinars,” *Physical Review D*, vol. 82, no. 4, p. 044009, 2010.
- [276] J. Keir, “Slowly decaying waves on spherically symmetric spacetimes and ultracompact neutron stars,” *Classical and Quantum Gravity*, vol. 33, no. 13, p. 135009, 2016.
- [277] M. Casals and A. Ottewill, “High-order tail in Schwarzschild spacetime,” *Physical Review D*, vol. 92, no. 12, p. 124055, 2015.
- [278] M. Calvani, F. de Felice, and L. Nobili, “Photon trajectories in the Kerr–Newman metric,” *Journal of Physics A*, vol. 13, no. 10, p. 3213, 1980.
- [279] Z. Stuchlík, “Null geodesics in the Kerr–Newman metric,” *Bulletin of the Astronomical Institutes of Czechoslovakia*, vol. 32, p. 366, 1981.
- [280] D. Pugliese, H. Quevedo, and R. Ruffini, “Equatorial circular orbits of neutral test particles in the Kerr–Newman spacetime,” *Physical Review D*, vol. 88, no. 2, p. 024042, 2013.
- [281] V. I. Dokuchaev, “Is there life inside black holes?,” *Classical and Quantum Gravity*, vol. 28, no. 23, p. 235015, 2011.

- [282] S. Ulbricht and R. Meinel, “A note on circular geodesics in the equatorial plane of an extreme Kerr–Newman black hole,” *Classical and Quantum Gravity*, vol. 32, no. 14, p. 147001, 2015.
- [283] F. S. Khoo and Y. C. Ong, “Lux in obscuro: photon orbits of extremal black holes revisited,” *Classical and Quantum Gravity*, vol. 33, no. 23, p. 235002, 2016. (Corrigendum: *Classical and Quantum Gravity*, vol. 34, no. 21, p. 219501, 2017).
- [284] Z. Stuchlík and S. Hledík, “Properties of the Reissner-Nordström spacetimes with a nonzero cosmological constant,” *Acta Physica Slovaca*, vol. 52, no. 5, p. 363, 2002.
- [285] A. Grenzebach, V. Perlick, and C. Lämmerzahl, “Photon regions and shadows of Kerr-Newman-NUT black holes with a cosmological constant,” *Physical Review D*, vol. 89, no. 12, p. 124004, 2014.
- [286] T. Igata, H. Ishihara, and Y. Takamori, “Stable bound orbits of massless particles around a black ring,” *Physical Review D*, vol. 87, no. 10, p. 104005, 2013.
- [287] T. Igata, “Stable bound orbits in six-dimensional Myers-Perry black holes,” *Physical Review D*, vol. 92, no. 2, p. 024002, 2015.
- [288] A. Wunsch, T. Müller, D. Weiskopf, and G. Wunner, “Circular orbits in the extreme Reissner-Nordström dihole metric,” *Physical Review D*, vol. 87, no. 2, p. 024007, 2013.
- [289] S. R. Dolan and J. O. Shipley, “Stable photon orbits in stationary axisymmetric electrovacuum spacetimes,” *Physical Review D*, vol. 94, no. 4, p. 044038, 2016.
- [290] P. O. Mazur, “Proof of uniqueness of the Kerr-Newman black hole solution,” *Journal of Physics A*, vol. 15, no. 10, p. 3173, 1982.
- [291] B. Carter, “Global structure of the Kerr family of gravitational fields,” *Physical Review*, vol. 174, no. 5, p. 1559, 1968.
- [292] G. P. Perry and F. I. Cooperstock, “Electrostatic equilibrium of two spherical charged masses in general relativity,” *Classical and Quantum Gravity*, vol. 14, no. 5, p. 1329, 1997.
- [293] G. G. Varzugin and A. S. Chistyakov, “Charged rotating black holes in equilibrium,” *Classical and Quantum Gravity*, vol. 19, no. 17, p. 4553, 2002.
- [294] G. A. Alekseev and V. A. Belinski, “Equilibrium configurations of two charged masses in general relativity,” *Physical Review D*, vol. 76, no. 2, p. 021501, 2007.
- [295] V. S. Manko, “Double-Reissner-Nordström solution and the interaction force between two spherical charged masses in general relativity,” *Physical Review D*, vol. 76, no. 12, p. 124032, 2007.
- [296] E. E. Zotos, “Classifying orbits in the classical Hénon–Heiles Hamiltonian system,” *Nonlinear Dynamics*, vol. 79, no. 3, p. 1665, 2015.

- [297] J. Grover and A. Wittig, “Black Hole Shadows and Invariant Phase Space Structures,” *Physical Review D*, vol. 96, no. 2, p. 024045, 2017.
- [298] J. Jia, X. Pang, and N. Yang, “Existence and stability of circular orbits in static and axisymmetric spacetimes,” *General Relativity and Gravitation*, vol. 50, no. 4, p. 41, 2018.
- [299] C. Charmousis, D. Langlois, D. Steer, and R. Zegers, “Rotating spacetimes with a cosmological constant,” *Journal of High Energy Physics*, vol. 2007, no. 02, p. 064, 2007.
- [300] M. Astorino, “Charging axisymmetric space-times with cosmological constant,” *Journal of High Energy Physics*, vol. 2012, no. 6, p. 86, 2012.
- [301] M. Astorino, “Stationary axisymmetric spacetimes with a conformally coupled scalar field,” *Physical Review D*, vol. 91, no. 6, p. 064066, 2015.
- [302] S. Hod, “On the number of light rings in curved spacetimes of ultra-compact objects,” *Physics Letters B*, vol. 776, p. 1, 2018.
- [303] H. Ishihara, M. Takahashi, and A. Tomimatsu, “Gravitational Faraday rotation induced by a Kerr black hole,” *Physical Review D*, vol. 38, no. 2, p. 472, 1988.
- [304] K. S. Thorne and R. D. Blandford, *Modern Classical Physics: Optics, Fluids, Plasmas, Elasticity, Relativity, and Statistical Physics*. Princeton University Press, 2017.
- [305] S. R. Dolan, “Scattering and absorption of gravitational plane waves by rotating black holes,” *Classical and Quantum Gravity*, vol. 25, no. 23, p. 235002, 2008.
- [306] L. C. S. Leite, S. R. Dolan, and L. C. B. Crispino, “Absorption of electromagnetic and gravitational waves by Kerr black holes,” *Physics Letters B*, vol. 774, p. 130, 2017.
- [307] A. I. Harte, “Gravitational lensing beyond geometric optics: I. Formalism and observables,” *General Relativity and Gravitation*, vol. 51, no. 1, p. 14, 2019.
- [308] A. I. Harte, “Gravitational lensing beyond geometric optics: II. Metric independence,” arXiv preprint arXiv:1906.10708 [gr-qc], 2019.
- [309] B. Mashhoon, “Electromagnetic scattering from a black hole and the glory effect,” *Physical Review D*, vol. 10, no. 4, p. 1059, 1974.
- [310] V. P. Frolov and A. A. Shoom, “Spinoptics in a stationary spacetime,” *Physical Review D*, vol. 84, no. 4, p. 044026, 2011.
- [311] V. P. Frolov and A. A. Shoom, “Scattering of circularly polarized light by a rotating black hole,” *Physical Review D*, vol. 86, no. 2, p. 024010, 2012.
- [312] T. Kashiwada, “On conformal Killing tensor,” *Natural Science Report, Ochanomizu University*, vol. 19, no. 2, p. 68, 1968.

- [313] S.-i. Tachibana, “On conformal Killing tensor in a Riemannian space,” *Tohoku Mathematical Journal, Second Series*, vol. 21, no. 1, p. 56, 1969.
- [314] B. Carter and R. G. McLenaghan, “Generalised master equations for wave equation separation in a Kerr or Kerr-Newman black hole background,” in *Proceedings of the second Marcel Grossmann meeting on general relativity* (R. Ruffini, ed.), p. 575, 1982.
- [315] B. Carter, “Black Hole Equilibrium States,” in *Black Holes* (C. DeWitt and B. S. DeWitt, eds.), p. 57, Gordon and Breach, 1973.
- [316] J. Ehlers and E. T. Newman, “The theory of caustics and wave front singularities with physical applications,” *Journal of Mathematical Physics*, vol. 41, no. 6, p. 3344, 2000.
- [317] S. Seitz, P. Schneider, and J. Ehlers, “Light propagation in arbitrary spacetimes and the gravitational lens approximation,” *Classical and Quantum Gravity*, vol. 11, no. 9, p. 2345, 1994.
- [318] H. Römer, “Theoretical optics,” *Wiley-VCH Verlag*, 2005.
- [319] A. Zenginoğlu and C. R. Galley, “Caustic echoes from a Schwarzschild black hole,” *Physical Review D*, vol. 86, no. 6, p. 064030, 2012.
- [320] S. R. Dolan, “Electromagnetic fields on Kerr spacetime, Hertz potentials and Lorenz gauge,” arXiv preprint arXiv:1906.04808 [gr-qc], 2019.
- [321] A. Wagemakers, A. Daza, and M. A. F. Sanjuán, “The saddle-straddle method to test for Wada basins,” arXiv preprint arXiv:1901.10728 [nlin.CD], 2019.
- [322] R. Barrio, F. Blesa, and S. Serrano, “Fractal structures in the Hénon-Heiles hamiltonian,” *EPL (Europhysics Letters)*, vol. 82, no. 1, p. 10003, 2008.
- [323] E. E. Zotos, “An overview of the escape dynamics in the Hénon–Heiles Hamiltonian system,” *Meccanica*, vol. 52, no. 11–12, p. 2615, 2017.
- [324] H. Goldstein, C. Poole, and J. Safko, *Classical Mechanics*. International Edition, Pearson Education International, 2002.
- [325] W. Kinnersley, “Type D vacuum metrics,” *Journal of Mathematical Physics*, vol. 10, no. 7, p. 1195, 1969.